

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE  
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

First Named  
Inventor : Anthony Aue et al.

Appln. No. : 10/813,208

Filed : March 30, 2004

For : STATISTICAL LANGUAGE MODEL FOR  
LOGICAL FORMS

Docket No. : M61.12-0630

Confirmation No.: 5138

Group Art Unit: 2626

Examiner: J. Stoffregen

---

**BRIEF FOR APPELLANTS**

*FILED ELECTRONICALLY AUGUST 21, 2008*

Sir:

This is an Appeal from an Office Action dated January 28, 2008 in which claims 1, 3-16 and 18-40 were finally rejected.

## **Contents**

---

<b>Real Party In Interest</b>	<b>3</b>
<b>No Related Appeals Or Interferences</b>	<b>3</b>
<b>Status Of The Claims</b>	<b>3</b>
<b>Status Of Amendments</b>	<b>3</b>
<b>Summary Of Claimed Subject Matter</b>	<b>4</b>
<b>I. Brief Background</b>	<b>4</b>
<b>II. The Claimed Subject Matter</b>	<b>5</b>
<b>A. Independent Claim 1 and Separately Argued Dependent             Claims 3, 4, 6, 9, 15, 18, 39 and 40</b>	<b>5</b>
<b>B. Independent Claim 23 and Separately Argued Dependent             Claim 24</b>	<b>7</b>
<b>C. Independent Claim 30</b>	<b>7</b>
<b>Grounds of Rejection To Be Reviewed On Appeal</b>	<b>8</b>
<b>Argument</b>	<b>8</b>
<b>I. Rejection of Claims 1, 7, 8, 16, 18-23, 39 and 40 under 35 USC §102(b)</b>	<b>8</b>
<b>A. Claims 1, 7, 8, 16</b>	<b>9</b>
<b>B. Claims 18-22</b>	<b>10</b>
<b>C. Claim 39</b>	<b>11</b>
<b>D. Claim 40</b>	<b>11</b>
<b>E. Claim 23</b>	<b>12</b>
<b>II. Rejection of 3-6, 9-15 and 24-38 under 35 USC §103(a)</b>	<b>12</b>
<b>A. Claim 3</b>	<b>13</b>
<b>B. Claims 4-5</b>	<b>13</b>
<b>C. Claim 6</b>	<b>14</b>
<b>D. Claims 9-14</b>	<b>14</b>
<b>E. Claim 15</b>	<b>14</b>
<b>F. Claims 24-29</b>	<b>15</b>
<b>G. Claims 30-38</b>	<b>15</b>
<b>Conclusion</b>	<b>16</b>
<b>Appendix A: Claims On Appeal</b>	<b>17</b>
<b>Appendix B: Cited References</b>	<b>25</b>
<b>Appendix C: Evidence Appendix</b>	<b>26</b>
<b>Appendix D: Related Proceedings Appendix</b>	<b>27</b>

**REAL PARTY IN INTEREST**

Microsoft Corporation, a corporation organized under the laws of the state of Washington, and having offices at One Microsoft Way, Redmond, Washington 98052, has acquired the entire right, title and interest in and to the invention, the application, and any and all patents to be obtained therefor, as set forth in the Assignment filed with the patent application and recorded on Reel 015163, frame 0222.

**NO RELATED APPEALS OR INTERFERENCES**

There are no known related appeals or interferences that will directly affect or be directly affected by or have a bearing on the Board's decision in this appeal.

**STATUS OF THE CLAIMS**

- I. Total number of claims in the application.
  - Claims in the application are: 1, 3-16 and 18-40
- II. Status of all the claims.
  - A. Claims canceled: 2 and 17
  - B. Claims withdrawn but not cancelled:
  - C. Claims pending: 1, 3-16 and 18-40
  - D. Claims allowed: none
  - E. Claims rejected: 1, 3-16 and 18-40
  - F. Claims Objected to: none
- III. Claims on appeal
  - The claims on appeal are: 1, 3-16 and 18-40

**STATUS OF AMENDMENTS**

No amendments have been filed after the final rejection.

## **SUMMARY OF CLAIMED SUBJECT MATTER**

### **I. Brief Background**

Machine translation systems receive a source language, translate the source language into a target language and provide an output in the target language. One example of a machine translation system uses logical forms or dependency graphs to translate. In this type of machine translation, a string in a source language is parsed to produce a source language logical form that is to be converted in to a target language logical form. A database of mappings include source language logical form pieces mapped to target language logical form pieces along with other metadata, such as sizes of mappings and frequencies of mapping are used for the conversion. Typically, the source language logical form piece of a mapping does not cover the entire source language logical form that was built from the source string. Therefore, a set of mappings must be selected and their target language logical form pieces combined to form a target logical form. This type of machine translation system first sorts by size, frequency and etc. to determine how well the source logical form pieces match the source logical form. Then, the sorted list is traversed in a top-down manner and the first set of compatible mappings is chosen. This system, however, fails test for all possible combinations of mapping, but just selects the first set of mappings that completely cover the source logical form.

Other types of machine translation systems do not employ logical forms and instead predict the most likely target language string given an input string in the source language using statistical models. These string-based statistical machine translation systems require a large search space, which means the system takes to long to consider all possible strings. Oftentimes to address this problem, the string-based machine translation systems make simplifying assumptions. Still other types of statistical machine translation systems map a syntax structure in the source language to a string in the target language. Oftentimes the events that depend on each other are closer together in a syntax tree that they are in the surface string of string-base models, however, distance between interdependent words can still be too large and similar concepts are expressed by different structures resulting in poor translation performance.



## **II. The Claimed Subject Matter**

### **A. Independent Claim 1 and Separately Argued Dependent Claims 3, 4, 6, 9, 15, 18, 39 and 40**

Independent claim 1 is directed to a method of decoding an input semantic structure (552, 1400) to generate an output semantic structure (556) (see FIG. 13, page 37, line 23 through page 38, line 6). The method includes providing a set of transfer mappings (542, FIGS. 15-21) that cover at least portions of the input semantic structure (552, 1400) (see page 38, lines 3-6). Each transfer mapping (542, FIGS. 15-21) has an input semantic side that describes nodes of the input semantic structure (552, 1400) and has an output semantic side that describes nodes of the output semantic structure (556) (see page 37, lines 14-22). A score for each of the set of transfer mappings (542, FIGS. 15-21) that describe a select node of the input semantic structure (552, 1400) is calculated (1312) using a statistical model (see page 26, line 1 through page 28, line 9). Which of the transfer mappings (542, FIGS. 15-21) that describes the select node and has a highest score is selected (1316) (see page 41, line 11 through page 42, line 3). The selected transfer mapping is used to construct the output semantic structure (556) (see page 47, lines 22-29).

Dependent claim 3 depends from independent claim 1 and further describes that when the score is calculated (1312) the score is calculated using a target language model that provides a probability of a set of nodes appearing in the output semantic structure (see page 29, line 1 through page 30, line 15 and page 44, line 1-10).

Dependent claim 4 depends from independent claim 1 and further describes that when the score is calculated (1312) the score is calculated using a channel model that provides a probability of an input semantic side of a transfer mapping given the output semantic side of the transfer mapping (see page 30, line 16 through page 34, line 16 and page 44, lines 11-15).

Dependent claim 6 depends from independent claim 1 and further describes that when the score is calculated (1312) the score is calculated using a fertility model that provides a probability of node deletion in a transfer mapping (see page 35, line 6 through page 36, line 15 and page 44, lines 16-19).

Dependent claim 9 depends from independent claim 1 and further describes that when the score is calculated (1312) of for each transfer mapping in the set of transfer mappings (542,

FIGS. 15-21) that separate scores are computed for a plurality of models and the separate scores are combined to determine the score for each transfer mapping (542, FIGS. 15-21) that describe a select node of the input semantic structure (552, 1400) (see page 44, line 1 through page 45, line 6).

Dependent claim 10 depends from dependent claim 9 and further describes the plurality of models as including a channel model that provides a probability of an input semantic side of a transfer mapping given the output semantic side of the transfer mapping (see page 30, line 16 through page 34, line 16 and page 44, lines 11-15).

Dependent claim 15 depends from dependent claim 9 and further describes multiplying each score by a weight to form weighted model scores and summing the weighted model scores to determine the score for each transfer mapping that describe a select node of the input semantic structure (552, 1400) (see page 14, line 27 through page 28, line 25).

Dependent claim 18 depends from dependent claim 16 and further describes that when a score is calculated for each of the set of transfer mappings (542, FIGS. 15-21) a score for a tree of transfer mappings is calculated by recursively calculating a score for each level of nested subtrees. Calculating a score for a subtree includes recursively scoring the subtrees of the subtree, calculating a score for the root transfer mapping of the subtree and combining the scores for the subtrees of the subtree with the score for the root transfer mapping of the subtree (see page 37, line 23 through page 45, line 10). A score for the root transfer mapping is calculated and the score for each subtree is combined with the score for the root transfer mapping (see page 45, line 11 through page 46, line 5).

Dependent claim 39 depends from independent claim 1 and further describes that when the set of transfer mappings are provided includes providing root transfer mappings that describe the select node and root transfer mappings that describe any child nodes of the select node (see page 38, lines 7-30).

Dependent claim 40 depends from dependent claim 29 and further describes that when the score is calculated for each of the set of transfer mappings, the calculation includes determining (1308) whether each transfer mapping that describes the select node also describes any child nodes (see page 38, line 21 through page 40, line 5). The score is calculated (1312) for each of the root transfer mappings that describe one of the child nodes of the select node with the

statistical model (see page 40, lines 6-29). Which of the root transfer mappings is selected (1316) that describe one of the child nodes of the select node that has the highest scores (see page 40, line 17 through page 42, line 3). The scores of the highest scoring root transfer mappings that describe each of the child nodes are combined with a score of the root transfer mapping of the select node to find the score for each of the set of transfer mappings that describes the select node (see page 42, line 4 through page 47, line 21).

#### **B. Independent Claim 23 and Separately Argued Dependent Claim 24**

Independent claim 23 is directed to a machine translation system (500) for translating an input in a first language into an output in a second language (see page 21, lines 15-26). The system (500) includes a parser (522) for parsing the input into an input semantic representation (552) (see page 24, lines 4-8). The system (500) includes a search component (524) configured to find a set of transfer mappings (542) (see page 24, lines 9-15). Each transfer mapping includes an input semantic side that corresponds with portions of the input semantic representation (552) (see page 37, lines 14-22). The system (500) includes a decoding component (554) configured to score each of the set of transfer mappings (542) that corresponds with a select portion of the input semantic representation (552) and to select which of the transfer mappings that correspond with the select portion of the input semantic representation (552) that has a highest score (see page 24, lines 16-25 and FIG. 13). The system (500) also includes a generation component (528) configured to generate the output (558) based on the selected transfer mapping (see page 25, lines 13-27).

Dependent claim 24 depends from claim 23 and further describes that the decoding component (554) scores each transfer mapping by using a plurality of statistical models (see page 26, line 1 through page 28, line 9).

#### **C. Independent Claim 30**

Independent claim 30 is directed to a method of determining a score for a word string. An input semantic structure (552, 1400) having a plurality of nodes that relate to an input word string is computed (see page 24, lines 4-8). A set of transfer mappings (542, FIGS. 15-21) are

obtained (see page 24, lines 9-15). Each of the set of transfer mappings including an input semantic side that describes at least one node of the input semantic structure (552, 1400) (see page 37, lines 14-22). Each of the set of transfer mappings (542, FIGS. 15-21) that describe a select node of the input semantic structure (552, 1400) are scored with a target language model that provides a probability of sequences of nodes appearing in an output semantic structure (556) having a plurality of nodes that relate to an output word string (558) (page 37, line 23 through page 47, line 29).

#### **GROUND OF REJECTION TO BE REVIEWED ON APPEAL**

- I. Whether claims 1, 7, 8, 16, 18-23, 39 and 40 are anticipated by Menezes et al. (see Appendix B, Exhibit A).
- II. Whether claims 3-6, 9-15 and 24-38 are obvious over Menezes et al. in view of Brown et al. (See Appendix B, Exhibit B).

#### **ARGUMENT**

##### **I. REJECTION OF CLAIMS 1, 7, 8, 16, 18-23, 39 and 40 UNDER 35 USC §102(b)**

Claims 1, 7, 8, 16, 18-23, 39, and 40 were rejected under 35 U.S.C. 102(b) as being anticipated by Menezes et al. (Appendix B, Exhibit A). "A claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference." *Verdegaal Bros. V. Union Oil Co. of California*, 2 USPQ2d 1051, 1053 (Fed. Circ. 1987).

Appellant respectfully believes that the Examiner has failed to present a case of anticipation against the rejected claims. Specifically, the cited reference fails to describe all of the elements in claims 1, 7, 8, 16, 18-23, 39, and 40.

Before discussing in detail the rejection under 35 U.S.C. 102(b), It is pointed out that the background of Appellant's patent application discusses that features of Menezes et al. reference are known. However, it is further pointed out that the Appellant's claims differentiate from that which was disclosed in the background. Like Appellant's patent application, Menezes et al. uses source logical forms as a basis for machine translation. To translate, however, Menezes et al.

uses transfer logical form mappings that are sorted by size, frequency and etc. Then the sorted mappings are analyzed to find the first set of compatible mappings that covers the source logical form. This type of system does not test all combinations of mappings as is discussed in Appellant's application, but only select the first set of mappings that completely cover the source logical form.

**A. Claims 1, 7, 8 and 16**

In regards to independent claim 1, Menezes et al. fails to describe "calculating a score for each of the set of transfer mappings that describe a select node of the input semantic structure using a statistical model." Instead, Menezes et al. describes (in paragraph 66) that "matching component 224 searches for the best set of matching transfer mappings in database 218 that have matching lemmas, parts of speech and other feature information." Menezes further describes in paragraph 66 and also lays out in paragraph 120 that "the set of best matches is found based on a predetermined metric. For example, transfer mappings having larger (more specific) logical forms may illustratively be preferred to transfer mappings having smaller (more general) logical forms. Among mappings having logical forms of equal size, matching component 224 may illustratively prefer higher frequency mappings. Mappings may also match overlapping portions of the source logical form 252 provided that they do not conflict with each other in any way. A set of mappings collectively may be illustratively preferred if they cover more of the input sentence than alternative sets." While Menezes et al. finds best matching transfer mappings for a source logical form, claim 1 actually calculates a score for each transfer mapping that describes a select node of the input semantic structure. By calculating a score, all transfer mappings that cover one or more nodes of the source logical form are considered. The Examiner considers the metrics illustrated on page 9, table of 1 of Menezes et al. as used for calculating a score. However, these metrics are used when finding best transfer mappings for an entire source logical form. There is no indication that these metrics are used to score transfer mappings for a select node.

In addition, Menezes et al. fails to describe "selecting which of the transfer mappings that describe the select node has a highest score" as claimed in claim 1. As previously discussed,

Menezes et al. determines the best transfer mappings for a source logical form. Menezes et al. fails to determine which score of a transfer mapping calculated for a select node is the highest score. The Examiner considers paragraph 121 of Menezes et al. as describing the selecting of a highest score of the scored transfer mapping for a select node. However, paragraph 121 merely states that “a subset of matching transfer mappings is selected.” There is no indication that a highest score let alone a highest score of transfer mappings for a select node is selected.

Still further, Menezes et al. fails to describe “using the selected transfer mapping to construct the output semantic structure” as claimed in claim 1. The Examiner considers paragraph 121 of Menezes et al. as describing this. However, paragraph 121 merely states that “transfer mappings in the subset are combined into a transfer logical form from which the output text is generated.” There is no indication that the selected transfer mapping that has a highest score that describes a select node is used to construct an output semantic structure.

It is respectfully submitted that claim 1 is in condition for allowance. It is respectfully submitted that claims 7, 8 and 16 are also in condition for allowance as depending on an allowable base claim.

#### **B. Claims 18- 22**

It is respectfully submitted that claim 18 is also in condition for allowance at least based on its dependence on allowable claims 16 and 1. However, claim 18 is in condition for allowance for additional reasons. In particular, Menezes et al. fails to describe calculating a score for each level of nested subtrees – where each subtree is calculated by calculating a score for the root transfer mapping of a subtree and combining the score of the subtrees of the subtree with the score of the root transfer mapping of the subtree. Menezes et al fails to describe that after a score is calculated for each level of nested subtrees, that a score is calculated for the root transfer mapping and then the root transfer mapping score is combined with each score of the subtree.

Beside Menezes et al. failing to calculate a score as noted above. Menezes et al. fails to calculate a score as is claimed in claim 18. The Examiner points to paragraphs 92 and 93 and table 1 of Menezes et al. as teaching the features of claim 18. However, paragraphs 92 and 93 discuss the alignment of logical forms for the training and building of a transfer mapping

database. The alignment of logical forms to make the transfer mapping database is much different than the scoring of select transfer mappings in the approach as claimed in claim 18 to find highest scoring mappings.

It is respectfully submitted that claim 18 is in condition for allowance. In addition, claims 19 -22 are in condition for allowance at least based on their dependence on allowable claim 18.

**C. Claim 39**

It is respectfully submitted that claim 39 is also in condition for allowance at least based on its dependence on allowable claim 1. However, claim 39 is in condition for allowance for additional reasons. In particular, Menezes et al. fails to describe providing transfer mappings for the select node as well as any child nodes of the select nodes as claimed in claim 39. The Examiner points to FIG. 5B to show the features of claim 39. However, FIG. 5B illustrates the alignment of a logical form diagramming a Spanish sentence to a logical form of the same sentence in English. Menezes fails to describe the use of multiple transfer mappings for each parent node and child nodes.

It is respectfully submitted that claim 39 is in condition for allowance.

**D. Claim 40**

It is respectfully submitted that claim 40 is also in condition for allowance at least based on its dependence on allowable claims 1 and 39. However, claim 40 is in condition for allowance for additional reasons. In particular, Menezes et al. fails to describe combining scores of the highest scoring mappings that describe the child node with a score of the select node to find the scores for each set of transfer mappings that describe the select node. Again, although Menezes et al. discusses alignment of logical forms for the training and building of a transfer mapping database, the alignment of logical forms to make the transfer mapping database is much different than the scoring of select transfer mappings as claimed in claim 40.

It is respectfully submitted that claim 40 is in condition for allowance

**E. Claim 23**

In regards to independent claim 23, it is respectfully submitted that Menezes et al. fails to describe “a decoding component configured to score each of the set of transfer mappings that corresponds with a select portion of the input semantic representation and to select which of the transfer mappings that correspond with the select portion of the input semantic representation has a highest score.” While Menezes et al. finds best matching transfer mappings for a source logical form, claim 23 actually uses a decoding component that scores for each of the set of transfer mappings that corresponds to a select portion of the input semantic representation. The Examiner considers the metrics discussed in paragraph 7 of Menezes et al. as describing such a decoding component. However, these metrics are used when finding best transfer mappings for an entire source logical form. There is no indication that these metrics are used to score transfer mappings for a select portion of the input semantic representation.

It is respectfully submitted that claim 23 is in condition for allowance.

**II. REJECTION OF CLAIMS 3-6, 9-15 and 24-38 UNDER 35 USC §103(a)**

Claims 3-6, 9-15, and 24-38 were rejected under 35 U.S.C. 103(a) as being unpatentable over Menezes et al (Appendix B, Exhibit A) in view of Brown et al.(Appendix B, Exhibit B). As established by *Graham v. John Deere Co.*, 148 USPQ 149 (1968), adopted in *KSR v. International Co. v. Teleflex Inc.*, 82 USP2d 1385 (2007) and as recited in MPEP §2141, obviousness is a question of law based on underlying factual inquiries. The factual inquiries enunciated by the Graham Court were as follows:

- (A) Ascertain the scope and content of the prior art;
- (B) Ascertain the differences between the claimed invention and the prior art; and
- (C) Resolve the level of ordinary skill in the pertinent art.

It is respectfully submitted that differences between the claimed invention and the combination of cited references still exist and therefore the Examiner as failed to clearly articulate reasons to support a legal conclusion of obviousness.

Before discussing in detail the rejection under 35 U.S.C. 103(a), It is pointed out that the background of Appellant’s patent application discusses that features of the Brown et al. reference



are known. However, it is further pointed out that the Appellant's claims differentiate from that which was disclosed in the background and differentiate from the combination of cited references.

**A. Claim 3**

It is respectfully submitted that claim 3 is also in condition for allowance at least based on its dependence on allowable claim 1. However, claim 3 is in condition for allowance for additional reasons. In particular, the combination of cited references et al. fail to describe where calculating a score for each transfer mapping that describes a select node includes calculating a score with a target language model. The Examiner points to col. 8, lines 48-50 of Brown et al. as teaching a probability model. However, Brown et al. fails to describe that such a probability model is used to calculate a score of a transfer mapping as is claimed. More specifically, the combination of references fail to teach how a probability model in machine translation could be used with a logical form machine translation system in the way in which is claimed.

It is respectfully submitted that claim 3 is in condition for allowance.

**B. Claims 4-5**

It is respectfully submitted that claim 4 is also in condition for allowance at least based on its dependence on allowable claim 1. However, claim 4 is in condition for allowance for additional reasons. In particular, the combination of cited references fail to describe where calculating a score for each transfer mapping that describes a select node includes calculating a score with a channel model. The Examiner points to col. 8, lines 52-54 of Brown et al. as teaching such a probability model. However, Brown et al. fails to describe that such a probability model is used to calculate a score of a transfer mapping as is claimed. More specifically, the combination of references fail to teach how a probability model in machine translation could be used with a logical form machine translation system in the way in which is claimed.

It is respectfully submitted that claim 4 is in condition for allowance. In addition, claim 5 is in condition for allowance at least based on its dependence on allowable claim 4.

**C. Claim 6**

It is respectfully submitted that claim 4 is also in condition for allowance at least based on its dependence on allowable claim 1. However, claim 4 is in condition for allowance for additional reasons. In particular, the combination of cited references fail to describe where calculating a score for each transfer mapping that describes a select node includes calculating a score with a fertility model. The Examiner points to col. 52, lines 35-6 of Brown et al. as teaching such a probability model. However, Brown et al. fails to describe that such a probability model is used to calculate a score of a transfer mapping as is claimed. More specifically, the combination of references fail to teach how a probability model in machine translation could be used with a logical form machine translation system in the way in which is claimed.

It is respectfully submitted that claim 6 is in condition for allowance.

**D. Claim 9-14**

It is respectfully submitted that claim 9 is also in condition for allowance at least based on its dependence on allowable claim 1. However, claim 9 is in condition for allowance for additional reasons. In particular, the combination of cited references fail to describe where calculating a score for each transfer mapping that describes a select node includes computing separate scores for a plurality of models and combining the separate score to determine the score for each transfer mapping. While the abstract of Brown et al. recites that two models are used to score a hypotheses and the two model scores are combined into one, the combination of cited references fail to describe that scores are calculated for each transfer mapping of a select node from a plurality of models.

It is respectfully submitted that claim 9 is in condition for allowance. In addition, claims 10-14 are in condition for allowance at least based on their dependence on allowable claim 9.

**E. Claim 15**

It is respectfully submitted that claim 15 is also in condition for allowance at least based on its dependence on allowable claims 1 and 9. However, claim 15 is in condition for allowance for additional reasons. In particular, the combination of cited references fail to describe where

combining separate score includes multiplying each score by a weight to form weighted model scores and summing the weighted model score to determine the score for each transfer mapping that describes a select node. The Examiner points to equation 7 of Brown et al. However, all the Examiner says is that this equation includes  $\lambda$  without pointing to anywhere in the Brown et al. specification that describes  $\lambda$  as being a weight, let alone that a weight is applied each of plurality of models scores or that the weight is used to calculate a score for each transfer mapping that describes a select node.

It is respectfully submitted that claim 15 is in condition for allowance.

**F. Claim 24-29**

It is respectfully submitted that claim 24 is also in condition for allowance at least based on its dependence on allowable claim 23. However, claim 24 is in condition for allowance for additional reasons. In particular, the combination of cited references fail to describe where the decoding component scores each transfer mapping by using a plurality of statistical models. Although the abstract of Brown et al. describes the use of two different models, the combination of cited references fail to describe how one would use statistical model in a semantic representations type of translation or that you would score each transfer mapping that corresponds with a select portion of the representation using statistical models.

It is respectfully submitted that claim 24 is in condition for allowance. In addition, claims 25-29 are in condition for allowance at least based on their dependence on allowable claim 24.

**G. Claims 30-38**

In regards to claim 30, the combination of cited references fail to describe “scoring each of the set of transfer mappings that describe a select node of the input semantic structure with a target language model that provides a probability of sequences of nodes appearing in an output semantic structure having a plurality of nodes that relate to an output work string.” While the combination of cited references finds best matching transfer mappings for a source logical form and describes a target language model, claim 30 actually scores each of the set of transfer

mappings that describe a select node of the input semantic structure with a target language model. The combination of cited references fail to describe this.

It is respectfully submitted that claim 30 is in condition for allowance. It is respectfully submitted that claims 31-38 are also in condition for allowance as depending on an allowable base claim.

### **CONCLUSION**

For the reasons discussed above, Appellant respectfully submits that claims 1, 3-16 and 18-40 are neither described by the references cited by the Examiner nor has the Examiner presented a convincing line of reasoning as to why an artisan would have found the claimed invention to have been anticipated or obvious in light of the teachings of the reference. Thus, Appellant respectfully requests that the Board reverse the Examiner and find all pending claims 1, 3-16 and 18-40 allowable.

The Director is authorized to charge any fee deficiency required by this paper or credit any overpayment to Deposit Account No. 23-1123.

Respectfully submitted,

WESTMAN, CHAMPLIN & KELLY, P.A.

By: /Leanne Taveggia Farrell/  
Leanne Taveggia Farrell, Reg. No. 53,675  
900 Second Avenue South, Suite 1400  
Minneapolis, Minnesota 55402-3244  
Phone: (612) 334-3222  
Fax: (612) 334-3312

LTF/jmt

**Appendix A: Claims On Appeal**

1. A method of decoding an input semantic structure to generate an output semantic structure, the method comprising:

providing a set of transfer mappings that cover at least portions of the input semantic structure, each transfer mapping having an input semantic side that describes nodes of the input semantic structure and having an output semantic side that describes nodes of the output semantic structure;

calculating a score for each of the set of transfer mappings that describe a select node of the input semantic structure using a statistical model;

selecting which of the transfer mappings that describe the select node has a highest score; and

using the selected transfer mapping to construct the output semantic structure.

3. The method of claim 1, wherein calculating a score for each transfer mapping in the set of transfer mappings that describe a select node of the input semantic structure comprises calculating the score using a target language model that provides a probability of a set of nodes appearing in the output semantic structure.

4. The method of claim 1, wherein calculating a score for each transfer mapping in the set of transfer mappings that describe a select node of the input semantic structure comprises calculating the score using a channel model that provides a probability of an input semantic side of a transfer mapping given the output semantic side of the transfer mapping.

5. The method of claim 4, wherein calculating a score using the channel model comprises normalizing a channel model score based on a number of overlapping nodes between transfer mappings.

6. The method of claim 1, wherein calculating a score for each transfer mapping in the set of transfer mappings that describe a select node of the input semantic structure comprises calculating the score using a fertility model that provides a probability of node deletion in a transfer mapping.
7. The method of claim 1, wherein calculating a score for each transfer mapping in the set of transfer mappings that describe a select node of the input semantic structure comprises calculating a size score based on a number of nodes in the input semantic side of the transfer mapping.
8. The method of claim 1, wherein calculating a score for each transfer mapping in the set of transfer mappings that describe a select node of the input semantic structure comprises calculating a rank score based on a number of matching binary features in the input semantic structure and the input semantic side of the transfer mapping.
9. The method of claim 1, wherein calculating a score for each transfer mapping in the set of transfer mappings that describe a select node of the input semantic structure comprises:
  - computing separate scores for a plurality of models; and
  - combining the separate scores to determine the score for each transfer mapping that describe a select node of the input semantic structure.
10. The method of claim 9 wherein the plurality of models comprises a channel model that provides a probability of an input semantic side of a transfer mapping given the output semantic side of the transfer mapping.
11. The method of claim 9 wherein the plurality of models comprises a fertility model that provides a probability of node deletion in a transfer mapping.

12. The method of claim 9 wherein the plurality of models comprises a target language model that provides a probability of a set of nodes appearing in the output semantic structure.
13. The method of claim 9 and further comprising:  
computing a size score for each transfer mapping that describes a select node of the input semantic structure, the size score based on a number of nodes in the input semantic side of each transfer mapping; and  
combining the size score with the separate scores for the plurality of models to determine the score for each transfer mapping that describe a select node of the input semantic structure.
14. The method of claim 9 and further comprising:  
computing a rank score for each transfer mapping that describe a select node of the input semantic structure, the rank score based on a number of matching binary features in the input semantic structure and the input semantic side of each transfer mapping; and  
combining the rank score with the separate scores for the plurality of models to determine the score for each transfer mapping that describe a select node of the input semantic structure.
15. The method of claim 9 wherein combining the separate scores comprises:  
multiplying each score by a weight to form weighted model scores; and  
summing the weighted model scores to determine the score for each transfer mapping that describe a select node of the input semantic structure.

16. The method of claim 1, wherein providing a set of transfer mappings comprises providing a set of transfer mappings arranged as a tree structure and multiple levels of nested subtrees comprising a root transfer mapping and subtrees, each subtree comprising a root transfer mapping, wherein each transfer mapping in the set of transfer mappings appears as a root transfer mapping in at least one of the tree and subtrees.

18. The method of claim 16 wherein calculating a score for each of the set of transfer mappings comprises calculating a score for a tree of transfer mappings through steps comprises:

    recursively calculating a score for each level of nested subtrees, wherein calculating a score for a subtree comprises recursively scoring the subtrees of the subtree, calculating a score for the root transfer mapping of the subtree, and combining the scores for the subtrees of the subtree with the score for the root transfer mapping of the subtree;  
    calculating a score for the root transfer mapping; and  
    combining the score for each subtree with the score for the root transfer mapping.

19. The method of claim 18 wherein computing a score for a root transfer mapping comprises computing a size score for the root transfer mapping based on a number of nodes in the input semantic side of the root transfer mapping.

20. The method of claim 18, wherein combining the score of subtrees with the score for a root transfer mapping comprises combining size scores for the subtrees with the size score for the root transfer mapping by averaging the size scores for the subtrees with the size score for the root transfer mapping.

21. The method of claim 18 wherein computing a score for a root transfer mapping comprises computing a rank score for the root transfer mapping based on a number of matching binary features in the input semantic structure and the input semantic side of the root transfer mapping.



22. The method of claim 21, wherein combining the score of subtrees with the score for a root transfer mapping comprises combining rank scores for the subtrees with the rank score of the root transfer mapping by averaging the rank scores for the subtrees with the rank score of the root transfer mapping.

23. A machine translation system for translating an input in a first language into an output in a second language, the system comprising:

- a parser for parsing the input into an input semantic representation;
- a search component configured to find a set of transfer mappings, wherein each transfer mapping includes an input semantic side that corresponds with portions of the input semantic representation;
- a decoding component configured to score each of the set of transfer mappings that corresponds with a select portion of the input semantic representation and to select which of the transfer mappings that correspond with the select portion of the input semantic representation has a highest score; and
- a generation component configured to generate the output based on the selected transfer mapping.

24. The machine translation system of claim 23, wherein the decoding component scores each transfer mapping by using a plurality of statistical models.

25. The machine translation system of claim 24, wherein the output comprises an output semantic representation and wherein the plurality of statistical models comprises a target model that provides a probability of a sequence of nodes appearing in the output semantic representation.

26. The machine translation system of claim 24, wherein the plurality of statistical models comprises a channel model that provides a probability of a set of semantic nodes in an input side of a transfer mapping given a set of semantic nodes in an output side of the transfer.

27. The machine translation system of claim 24, wherein the plurality of statistical models comprises a fertility model that provides a probability of a node deletion in the transfer mapping.

28. The machine translation system of claim 24, wherein the decoding component scores each transfer mapping using a size score based on a number of nodes in an input side of the transfer mapping.

29. The machine translation system of claim 24, wherein the decoding component scores each transfer mapping using a rank score based on a number of matching binary features between the input and an input side of the transfer mapping.

30. A method of determining a score for a word string, the method comprising:  
computing an input semantic structure having a plurality of nodes that relate to an input word string;  
obtaining a set of transfer mappings, each of the set of transfer mappings including an input semantic side that describes at least one node of the input semantic structure; and  
scoring each of the set of transfer mappings that describe a select node of the input semantic structure with a target language model that provides a probability of sequences of nodes appearing in an output semantic structure having a plurality of nodes that relate to an output word string.

31. The method of claim 30, wherein providing an input semantic structure having a plurality of nodes comprises providing an input semantic structure having a plurality of word nodes and at least one relationship node that describes a semantic relationship between words.

32. The method of claim 30, wherein providing word nodes comprises providing word nodes for lemmas.

33. The method of claim 30, wherein scoring the input word string with a target language model comprises scoring the input word string with the target language model in machine translation.

34. The method of claim 30, wherein scoring the input word string with a target language model comprises scoring the input word string with the target language model in speech recognition.

35. The method of claim 30, wherein scoring the input word string with a target language model comprises scoring the input word string with the target language model in optical character recognition.

36. The method of claim 30, wherein scoring the input word string with a target language model comprises scoring the input word string with the target language model in grammar checking.

37. The method of claim 30, wherein scoring the input word string with a target language model comprises scoring the input word string with the target language model in handwriting recognition.

38. The method of claim 30, wherein scoring the input word string with a target language model comprises scoring the input word string with the target language model in information extraction.

39. The method of claim 1, wherein providing the set of transfer mappings comprises providing root transfer mappings that describe the select node and root transfer mappings that describe any child nodes of the select node.

40. The method of claim 39, wherein calculating the score for each of the set of transfer mappings comprises:

determining whether each transfer mapping that describes the select node also describes any child nodes;

calculating a score for each of the root transfer mappings that describe one of the child nodes of the select node with the statistical model;

selecting which of the root transfer mappings that describe one of the child nodes of the select node have highest scores;

combining scores of the highest scoring root transfer mappings that describe each of the child nodes with a score of the root transfer mapping of the select node to find the score for each of the set of transfer mappings that describes the select node.

**Appendix B: Cited References**

Exhibit A – Menezes et al., U.S. Publication No. 2003/0023422, filed July 5, 2001

Exhibit B – Brown et al., U.S. Patent No. 5,477,451, filed July 25, 1991

**Appendix C: Evidence Appendix**

There is no known evidence submitted pursuant to 37 CFR §§ 1.130, 1.131 or 1.132 or other evidence entered by the Examiner.

**Appendix D: Related Proceedings Appendix**

There are no known related appeals or interferences regarding the present appeal.

# **EXHIBIT A**

U.S. Publication No. 2003/0023422

Filed July 5, 2001

Menezes et al.





US 2003/0023422A1

(19) **United States**(12) **Patent Application Publication** (110) Pub. No.: **US 2003/0023422 A1**

Menezes et al.

(43) Pub. Date:

**Jan. 30, 2003**(54) **SCALEABLE MACHINE TRANSLATION SYSTEM**

Related U.S. Application Data

(60) Provisional application No. 60/295,338, filed on Jun. 1, 2001.

(76) Inventors: Arul A. Menezes, Bellevue, WA (US);  
Stephen D. Richardson, Redmond, WA (US);  
Jessie E. Pinkham, Bellevue, WA (US);  
William R. Dolan, Redmond, WA (US)

Publication Classification

(51) Int. Cl.<sup>7</sup> G06F 17/28

(52) U.S. Cl. 704/2

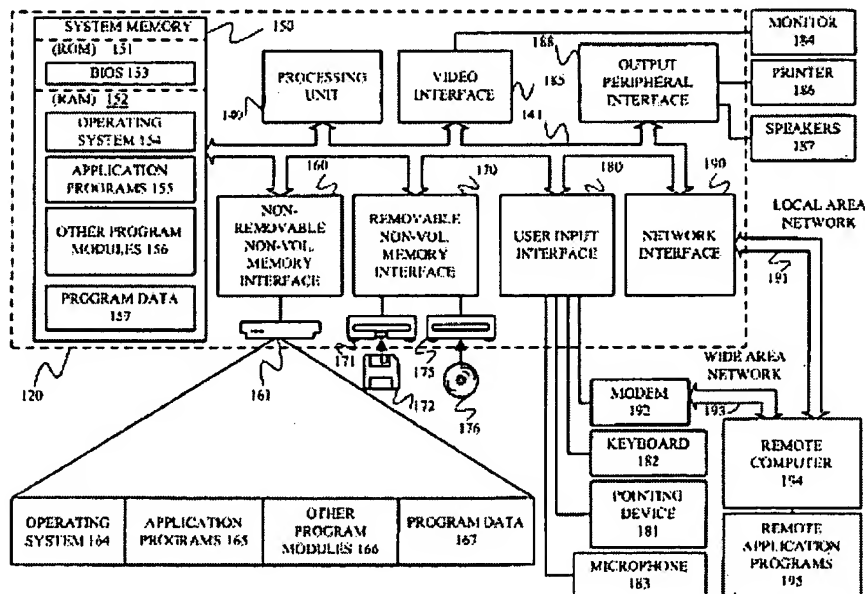
Correspondence Address:

Joseph R. Kelly  
WESTMAN CHAMPLIN & KELLY  
Suite 1600-International Centre  
900 South Second Avenue  
Minneapolis, MN 55402-3319 (US)(57) **ABSTRACT**

A computer implemented method translates a textual input in a first language to a textual output in a second language. An input logical form is generated based on the textual input. When a plurality of transfer mappings in a transfer mapping database match the input logical form (or at least a portion thereof) one or more of these plurality of matching transfer mappings is selected based on a predetermined metric. Textual output is generated based on the selected transfer logical form.

(21) Appl. No.: 09/899,755

(22) Filed: Jul. 5, 2001



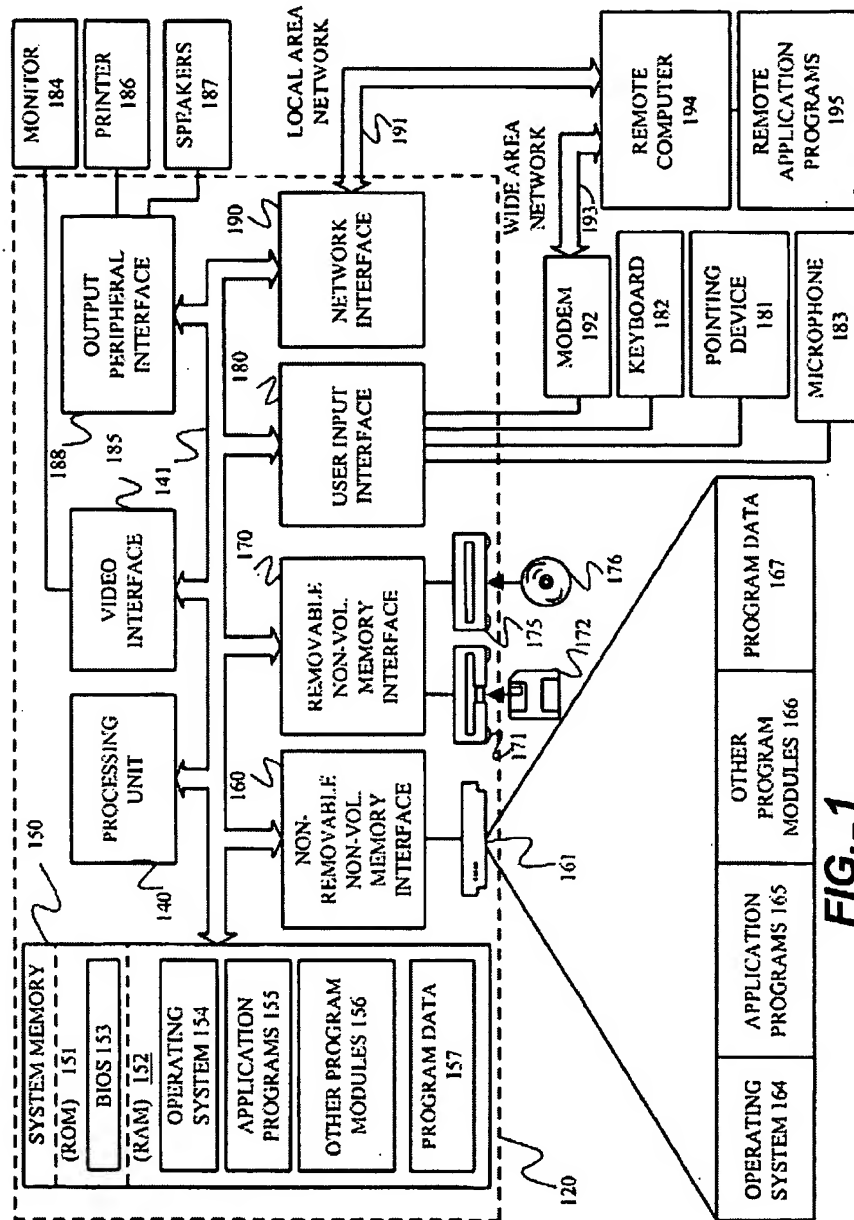


FIG. 1




FIG. 2B






FIG. 2

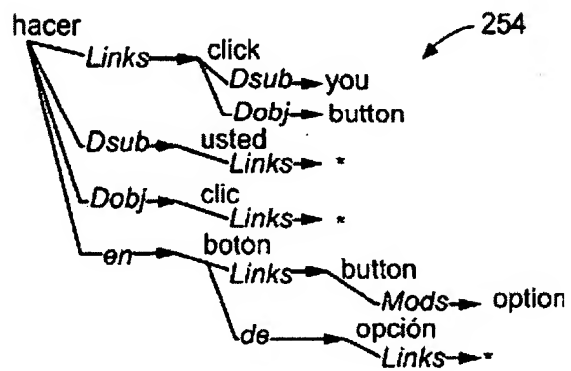
FIG. 2A




**FIG.\_3A**




**FIG.\_3B**




**FIG.\_3C**

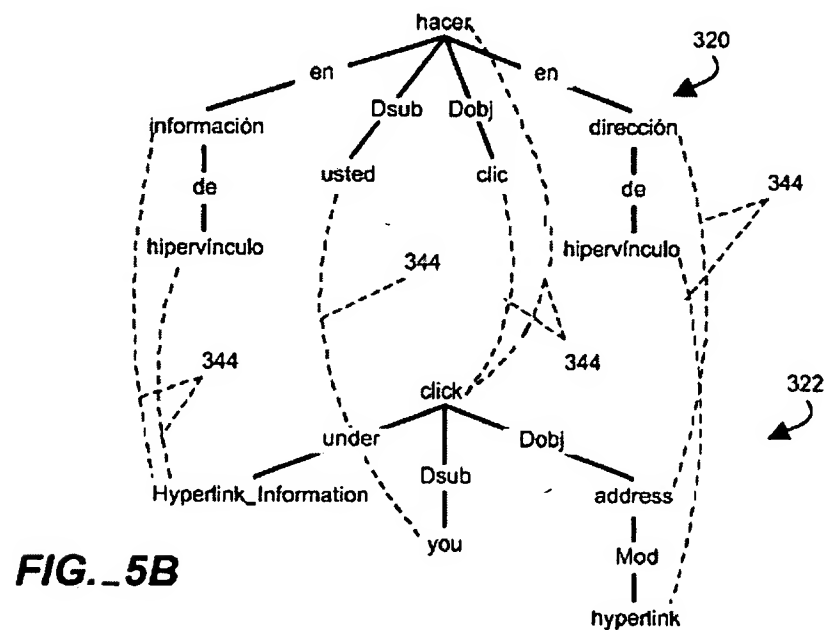
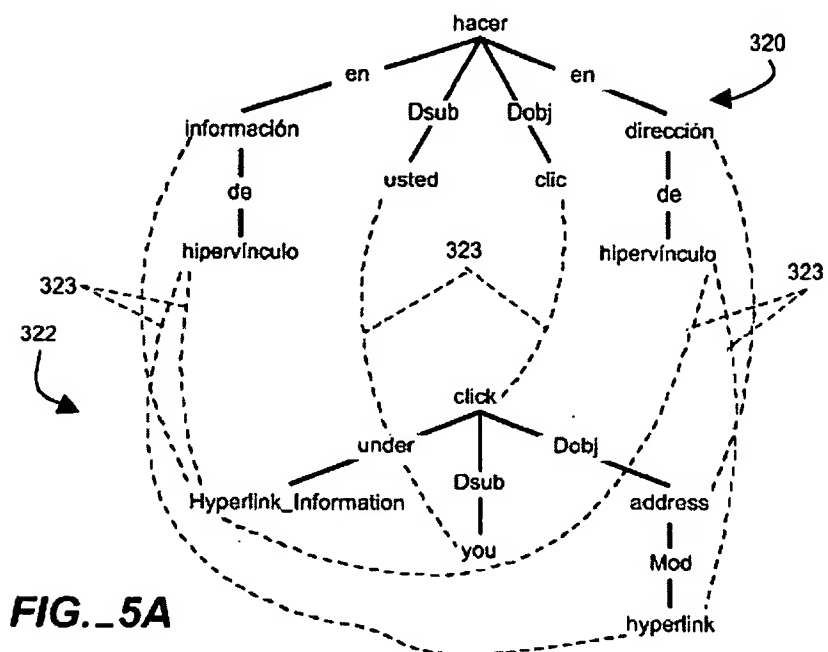


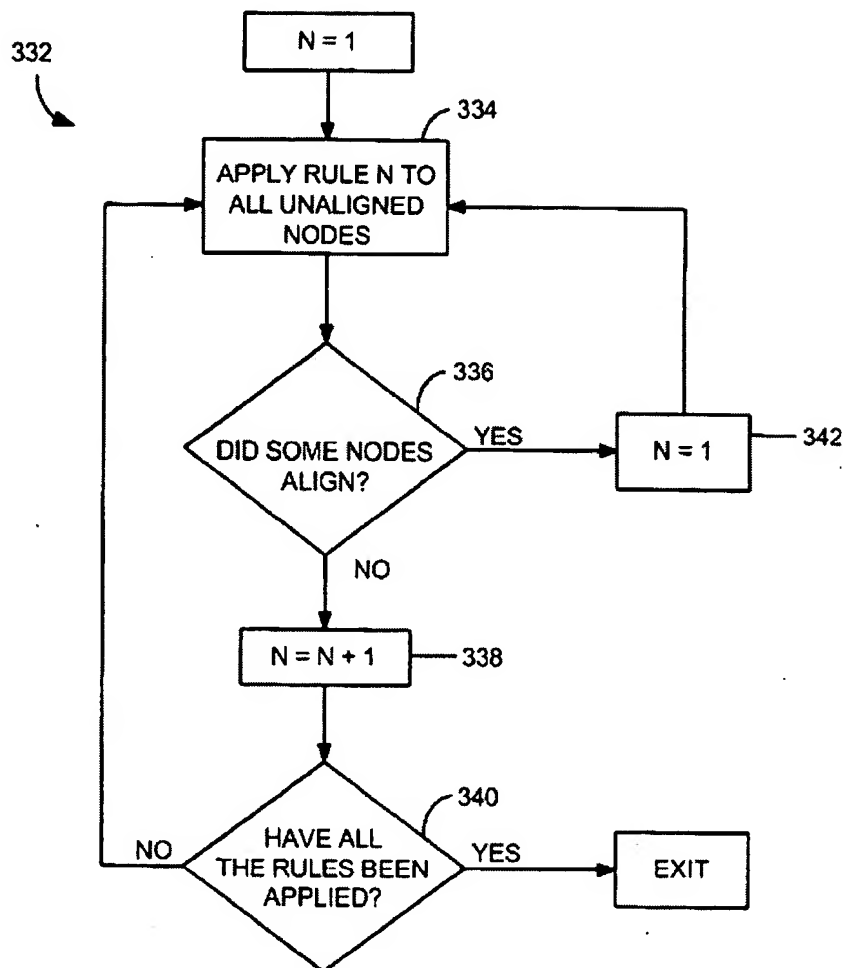
**FIG.\_4**



**FIG.\_6**

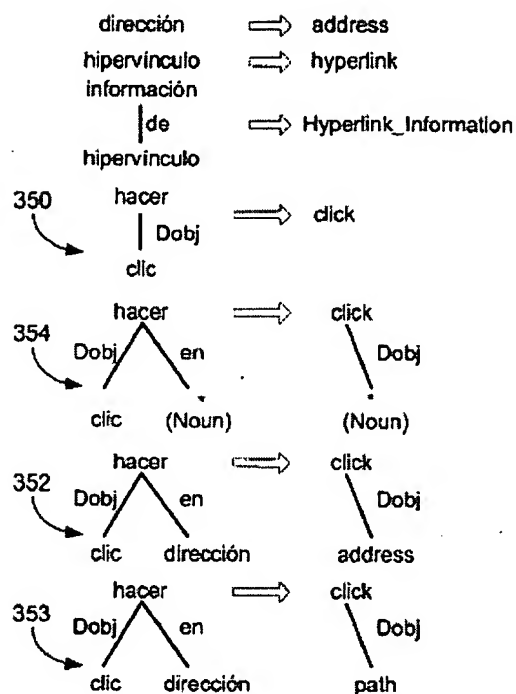




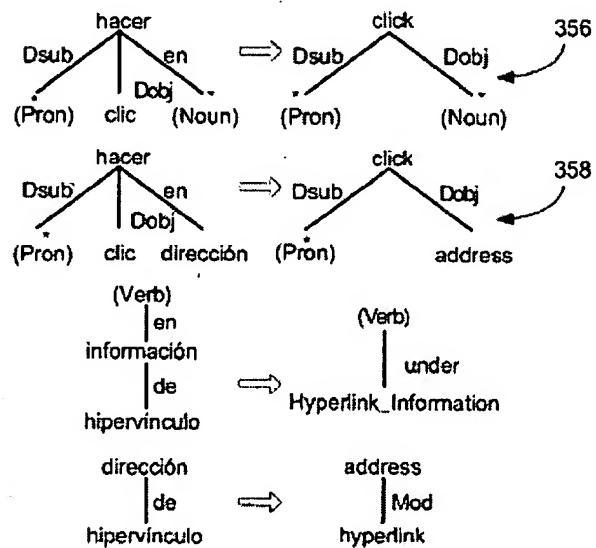


**FIG. 7**





**FIG.\_8**



# SCALABLE MACHINE TRANSLATION SYSTEM CROSS-REFERENCE TO RELATED APPLICATION

[0001] This application claims the benefit of U.S. provisional patent application serial No. 60/295,338, filed Jan. 1, 2001.

## BACKGROUND OF THE INVENTION

[0002] The present invention relates to automated language translation systems. More specifically, the present invention relates to a scalable machine translation system and architecture.

[0003] Machine translation systems are systems which receive a textual input in one language, translate it to a second language, and provide a textual output in the second language. Current commercially available machine translation systems rely on hand-coded transfer components that are both difficult and expensive to customize for a particular domain, and are also very difficult to scale to a desirable size. These disadvantages have limited their cost effectiveness and overall utility.

[0004] A variety of example based machine translation systems have been created to address these deficiencies. A number of such systems are described in H. Somers, *Review Article: Example-Based Machine Translation*, Machine Translation 14:113,157, 1999. Some of these typical example based machine translation research systems have been built with an example base built from up to approximately 200 sentences. They have encountered a great deal of difficulty in scaling to a larger example base and the performance of the system suffers from this difficulty.

[0005] Other of the data driven systems described in Somers parse the inputs from the example base using different parsers, based upon the particular language of the input text. The dependency structures resulting from such parsing are thus different, based upon the language and the particular parsing strategy used. Therefore, comparing the dependency structures from one language to the next is difficult, if not impossible.

[0006] Such prior systems have also not been easily scalable. For example, in order to increase the number of sentences over and above, for example, 200 sentences or so, has been very difficult. This is because the prior systems have difficulty handling noisy input data. Instead, the input data has been required to be in a precise form, or it has been cleaned up, and placed in the proper form, by hand. Of course, this makes it very difficult to dramatically increase the number of sentences because of the intensive labor required to clean up the data.

## SUMMARY OF THE INVENTION

[0007] A computer implemented method translates a textual input in a first language to a textual output in a second language. An input logical form is generated based on the textual input. When a plurality of transfer mappings in a transfer mapping database match the input logical form (or at least a portion thereof) one or more of those plurality of matching transfer mappings is selected based on a predetermined metric. These transfer mappings are stitched together to form a transfer logical form. The textual output is generated based on the transfer logical form.

[0008] A transfer mapping is illustratively composed of a pair of logical form fragments, including a source and target logical form (L.F.), learned from the training data. At runtime the source side of these mappings is matched against the input. Among such matched mappings, a set is chosen. The target sides of these mappings is then stitched together to produce a single target L.F. The output string is then generated from the target L.F.

[0009] The predetermined metric can take one of a variety of forms, including the number of input nodes covered by the set of mappings collectively, size of the different transfer mappings that match the input logical form, the frequency with which the plurality of matching transfer mappings were generated during a training phase used in training the transfer mapping database, frequencies with which the plurality of matching transfer mappings are generated from completely aligned logical forms during training, frequencies with which the plurality of matching transfer mappings were generated from non-fitted parses of the training data, and a score associated with each of the plurality of matching transfer mappings that is indicative of a confidence in the transfer mapping with which it is associated.

[0010] The present invention can also be embodied as a machine translation system including a matching component configured to implement the method discussed above.

[0011] The present invention can also be implemented as a machine translation system that includes an input generator generating an input dependency structure based on the textual input. The system also includes a transfer mapping database that holds a plurality of transfer mapping dependency structures formed based on at least 10,000 parallel, aligned, training sentences. The transfer mapping database can also be formed based on 50,000, 100,000, 180,000, or even in excess of 200,000 training sentences.

[0012] In addition, the present invention can be embodied as a computer implemented method of training a transfer mapping database which includes generating shared input logical forms for bilingual input sentences, the input logical forms being shared across both languages.

[0013] In yet another embodiment, the present invention trains the transfer mapping database by filtering transfer mappings obtained from aligned logical forms, aligned during training.

## BRIEF DESCRIPTION OF THE DRAWINGS

[0014] FIG. 1 is a block diagram of an illustrative environment in which the present invention may be used.

[0015] FIG. 2 is a block diagram of a machine translation architecture in accordance with one embodiment of the present invention.

[0016] FIG. 3A is an example of a logical form produced for a textual input in a source language (in this example, Spanish).

[0017] FIG. 3B is a linked logical form for the textual input in the source language.

[0018] FIG. 3C is a target logical form representing a translation of the source language input to a target language output (in this example, English).

[0019] FIG. 4 is a flow diagram illustrating a method for aligning nodes.

[0020] FIG. 5A is an example of tentative correspondences formed between logical forms.

[0021] FIG. 5B is an example of aligned nodes formed between the logical forms of FIG. 5A.

[0022] FIG. 6 is a flow diagram illustrating application of a set of rules to the method of FIG. 4.

[0023] FIG. 7 is a flow diagram illustrating application of an ordered set of rules.

[0024] FIG. 8 is a set of transferred mappings associated with the example of FIG. 5B.

#### DETAILED DESCRIPTION OF ILLUSTRATIVE EMBODIMENTS

##### General Overview

[0025] The following is a brief description of a general purpose computer 120 illustrated in FIG. 1. However, the computer 120 is only one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of the invention. Neither should the computer 120 be interpreted as having any dependency or requirement relating to any one or combination of modules illustrated therein.

[0026] The invention may be described in the general context of computer-executable instructions, such as program modules, being executed by a computer. Generally, program modules include routines, programs, objects, modules, data structures, etc. that perform particular tasks or implement particular abstract data types. The invention may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote computer storage media including memory storage devices. Tasks performed by the programs and modules are described below and with the aid of figures. Those skilled in the art can implement the description and figures as processor executable instructions, which can be written on any form of a computer readable media.

[0027] With reference to FIG. 1, modules of computer 120 may include, but are not limited to, a processing unit 140, a system memory 150, and a system bus 141 that couples various system modules or components including the system memory to the processing unit 140. The system bus 141 may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry Standard Architecture (ISA) bus, Universal Serial Bus (USB), Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Module Interconnect (PCI) bus also known as Mezzanine bus. Computer 120 typically includes a variety of computer readable mediums. Computer readable mediums can be any available media that can be accessed by computer 120 and includes both volatile and nonvolatile media, removable and non-

removable media. By way of example, and not limitation, computer readable mediums may comprise computer storage media and communication media. Computer storage media includes both volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program modules/components or other data. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical disk storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by computer 120.

[0028] Communication media typically embodies computer readable instructions, data structures, program modules or other data in a modulated data signal such as a carrier wave or other transport mechanism and includes any information delivery media. The term "modulated data signal" means a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communication media includes wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, FR, infrared and other wireless media. Combinations of any of the above should also be included within the scope of computer readable media.

[0029] The system memory 150 includes computer storage media in the form of volatile and/or nonvolatile memory such as read only memory (ROM) 151 and random access memory (RAM) 152. A basic input/output system 153 (BIOS), containing the basic routines that help to transfer information between elements within computer 120, such as during start-up, is typically stored in ROM 151. RAM 152 typically contains data and/or program modules that are immediately accessible to and/or presently being operated on by processing unit 140. By way of example, and not limitation, FIG. 1 illustrates operating system 154, application programs 155, other program modules 156, and program data 157.

[0030] The computer 120 may also include other removable/non-removable volatile/nonvolatile computer storage media. By way of example only, FIG. 1 illustrates a hard disk drive 161 that reads from or writes to non-removable, nonvolatile magnetic media, a magnetic disk drive 171 that reads from or writes to a removable, nonvolatile magnetic disk 172, and an optical disk drive 175 that reads from or writes to a removable, nonvolatile optical disk 176 such as a CD ROM or other optical media. Other removable/non-removable, volatile/nonvolatile computer storage media that can be used in the exemplary operating environment include, but are not limited to, magnetic tape cassettes, flash memory cards, digital versatile disks, digital video tape, solid state RAM, solid state ROM, and the like. The hard disk drive 161 is typically connected to the system bus 141 through a non-removable memory interface such as interface 160, and magnetic disk drive 171 and optical disk drive 175 are typically connected to the system bus 141 by a removable memory interface, such as interface 170.

[0031] The drives and their associated computer storage media discussed above and illustrated in FIG. 1, provide

storage of computer readable instructions, data structures, program modules and other data for the computer 120. In FIG. 1, for example, hard disk drive 161 is illustrated as storing operating system 164, application programs 165, other program modules 166, and program data 167. Note that these modules can either be the same as or different from operating system 154, application programs 155, other program modules 156, and program data 157. Operating system 164, application programs 165, other program modules 166, and program data 167 are given different numbers here to illustrate that, at a minimum, they are different copies.

[0032] A user may enter commands and information into the computer 120 through input devices such as a keyboard 182, a microphone 183, and a pointing device 181, such as a mouse, trackball or touch pad. Other input devices (not shown) may include a joystick, game pad, satellite dish, scanner, or the like. These and other input devices are often connected to the processing unit 140 through a user input interface 180 that is coupled to the system bus, but may be connected by other interface and bus structures, such as a parallel port, game port or a universal serial bus (USB). A monitor 184 or other type of display device is also connected to the system bus 141 via an interface, such as a video interface 185. In addition to the monitor, computers may also include other peripheral output devices such as speakers 187 and printer 186, which may be connected through an output peripheral interface 188.

[0033] The computer 120 may operate in a networked environment using logical connections to one or more remote computers, such as a remote computer 194. The remote computer 194 may be a personal computer, a hand-held device, a server, a router, a network PC, a peer device or other common network node, and typically includes many or all of the elements described above relative to the computer 120. The logical connections depicted in FIG. 1 include a local area network (LAN) 191 and a wide area network (WAN) 193, but may also include other networks. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets and the Internet.

[0034] When used in a LAN networking environment, the computer 120 is connected to the LAN 191 through a network interface or adapter 190. When used in a WAN networking environment, the computer 120 typically includes a modem 192 or other means for establishing communications over the WAN 193, such as the Internet. The modem 192, which may be internal or external, may be connected to the system bus 141 via the user input interface 180, or other appropriate mechanism. In a networked environment, program modules depicted relative to the computer 120, or portions thereof, may be stored in the remote memory storage device. By way of example, and not limitation, FIG. 1 illustrates remote application programs 195 as residing on remote computer 194. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

[0035] The invention is also operational with numerous other general purpose or special purpose computing systems, environments or configurations. Examples of well known computing systems, environments, and/or configura-

tions that may be suitable for use with the invention include, but are not limited to, regular telephones (without any screen) personal computers, server computers, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and the like.

#### Overview of Machine Translation System

[0036] Prior to discussing the present invention in greater detail, a brief discussion of a logical form may be helpful. A full and detailed discussion of logical forms and systems and methods for generating them can be found in U.S. Pat. No. 5,966,666 to Heidorn et al., issued Oct. 12, 1999 and entitled METHOD AND SYSTEM FOR COMPUTING SEMANTIC LOGICAL FORMS FROM SYNTAX TREES. Briefly, however, logical forms are generated by performing a morphological analysis on an input text to produce conventional phrase structure analyses augmented with grammatical relations. Syntactic analyses undergo further processing in order to derive logical forms, which are data structures that describe labeled dependencies among content words in the textual input. Logical forms can normalize certain syntactical alternations, (e.g., active/passive) and resolve both intrasentential anaphora and long distance dependencies. As illustrated herein, for example in FIG. 3A, a logical form 252 can be represented as a graph, which helps intuitively in understanding the elements of logical forms. However, as appreciated by those skilled in the art, when stored on a computer readable medium, the logical forms may not readily be understood as representing a graph.

[0037] Specifically, a logical relation consists of two words joined by a directional relation type, such as: LogicalSubject, LogicalObject,

[0038] IndirectObject;

[0039] LogicalNominative, LogicalComplement, LogicalAgent;

[0040] CoAgent, Beneficiary;

[0041] Modifier, Attribute, SentenceModifier;

[0042] PrepositionalRelationship;

[0043] Synonym, Equivalence, Apposition;

[0044] Hypernym, Classifier, SubClass;

[0045] Means, Purpose;

[0046] Operator, Modal, Aspect, DegreeModifier, Intensifier;

[0047] Focus, Topic;

[0048] Duration, Time;

[0049] Location, Property, Material, Manner, Measure, Color, Size;

[0050] Characteristic, Part;

[0051] Coordinate;

[0052] User, Possessor;

[0053] Source, Goal, Cause, Result; and

[0054] Domain.

[0055] A logical form is a data structure of connected logical relations representing a single textual input, such as a sentence or part thereof. The logical form minimally consists of one logical relation and portrays structural relationships (i.e., syntactic and semantic relationships), particularly argument and/or adjunct relation(s) between important words in an input string.

[0056] In one illustrative embodiment, the particular code that builds logical forms from syntactic analyses is shared across the various source and target languages that the machine translation system operates on. The shared architecture greatly simplifies the task of aligning logical form segments from different languages since superficially distinct constructions in two languages frequently collapse onto similar or identical logical form representations. Examples of logical forms in different languages are described in greater detail below with respect to FIGS. 3A-3C.

[0057] FIG. 2 is a block diagram of an architecture of a machine translation system 200 in accordance with one embodiment of the present invention. System 200 includes parsing components 204 and 206, statistical word association learning component 208, logical form alignment component 210, lexical knowledge base building component 212, bilingual dictionary 214, dictionary merging component 216, transfer mapping database 218 and updated bilingual dictionary 220. During training and translation run time, the system 200 utilizes analysis component 222, matching component 224, transfer component 226 and/or generation component 228.

[0058] In one illustrative embodiment, a bilingual corpus is used to train the system. The bilingual corpus includes aligned translated sentences (e.g., sentences in a source or target language, such as English, in 1-to-1 correspondence with their human-created translations in the other of the source or target language, such as Spanish). During training, sentences are provided from the aligned bilingual corpus into system 200 as source sentences 230 (the sentences to be translated), and as target sentences 232 (the translation of the source sentences). Parsing components 204 and 206 parse the sentences from the aligned bilingual corpus to produce source logical forms 234 and target logical forms 236.

[0059] During parsing, the words in the sentences are converted to normalized word forms (lemmas) and can be provided to statistical word association learning component 208. Both single word and multi-word associations are iteratively hypothesized and scored by learning component 208 until a reliable set of each is obtained. Statistical word association learning component 208 outputs learned single word translation pairs 238 as well as multi-word pairs 240.

[0060] The multi-word pairs 240 are provided to a dictionary merge component 216, which is used to add additional entries into bilingual dictionary 214 to form updated bilingual dictionary 220. The new entries are representative of the multi-word pairs 240.

[0061] The single and multi-word pairs 238, along with source logical forms 234 and target logical forms 236 are provided to logical form alignment component 210. Briefly, component 210 first establishes tentative correspondences between nodes in the source and target logical forms 230 and 236, respectively. This is done using translation pairs from a bilingual lexicon (e.g. bilingual dictionary) 214, which can

be augmented with the single and multi-word translation pairs 238, 240 from statistical word association learning component 208. After establishing possible correspondences, alignment component 210 aligns logical form nodes according to both lexical and structural considerations and creates word and/or logical form transfer mappings 242. This aspect will be explained in greater detail below.

[0062] Basically, alignment component 210 draws links between logical forms using the bilingual dictionary information 214 and single and multi-word pairs 238, 240. The transfer mappings are optionally filtered based on a frequency with which they are found in the source and target logical forms 234 and 236 and are provided to a lexical knowledge base building component 212.

[0063] While filtering is optional, in one example, if the transfer mapping is not seen at least twice in the training data, it is not used to build transfer mapping database 218, although any other desired frequency can be used as a filter as well. It should also be noted that other filtering techniques can be used as well, other than frequency of appearance. For example, transfer mappings can be filtered based upon whether they are formed from complete parses of the input sentences and based upon whether the logical forms used to create the transfer mappings are completely aligned.

[0064] Component 212 builds transfer mapping database 218, which contains transfer mappings that basically link words and/or logical forms in one language, to words and/or logical forms in the second language. With transfer mapping database 218 thus created, system 200 is now configured for runtime translations.

[0065] During translation run time, a source sentence 250, to be translated, is provided to analysis component 222. Analysis component 222 receives source sentence 250 and creates a source logical form 252 based upon the source sentence input. An example may be helpful. In the present example, source sentence 250 is a Spanish sentence "Haga click en el boton de opcion" which is translated into English as "Click the option button" or, literally, "Make click in the button of option".

[0066] FIG. 3A illustrates the source logical form 252 generated for source sentence 250 by analysis component 222. The source logical form 252 is provided to matching component 224. Matching component 224 attempts to match the source logical form 252 to logical forms in the transfer mapping database 218 in order to obtain a linked logical form 254. Multiple transfer mappings may match portions of source logical form 252. Matching component 224 searches for the best set of matching transfer mappings in database 218 that have matching lemmas, parts of speech, and other feature information. The set of best matches is found based on a predetermined metric. For example, transfer mappings having larger (more specific) logical forms may illustratively be preferred to transfer mappings having smaller (more general) logical forms. Among mappings having logical forms of equal size, matching component 224 may illustratively prefer higher frequency mappings. Mappings may also match overlapping portions of the source logical form 252 provided that they do not conflict with each other in any way. A set of mappings collectively may be illustratively preferred if they cover more of the input sentence than alternative sets. Other metrics used in matching the input logical form to those found in database 218 are discussed in greater detail below with respect to Table 1.

[0067] After a set of matching transfer mappings is found, matching component 224 creates links on nodes in the source logical form 252 to copies of the corresponding target words or logical form segments received by the transfer mappings, to generate linked logical form 254. FIG. 3B illustrates an example of linked logical form 254 for the present example. Links for multi-word mappings are represented by linking the root nodes (e.g., Hacer and Click) of the corresponding segments, then linking an asterisk to the other source nodes participating in the multi-word mapping (e.g., Used and Click). Sublinks between corresponding individual source and target nodes of such a mapping (not shown in FIG. 3B) may also illustratively be created for use during transfer.

[0068] Transfer component 226 receives linked logical form 254 from matching component 224 and creates a target logical form 256 that will form the basis of the target translation. This is done by performing a top down traversal of the linked logical form 254 in which the target logical form segments pointed to by links on the source logical form 252 nodes are combined. When combining together logical form segments for possibly complex multi-word mappings, the sublinks set by matching component 224 between individual nodes are used to determine correct attachment points for modifiers, etc. Default attachment points are used if needed.

[0069] In cases where no applicable transfer mappings are found, the nodes in source logical form 252 and their relations are simply copied into the target logical form 256. Default single word translations may still be found in transfer mapping database 218 for these nodes and inserted in target logical form 256. However, if none are found, translations can illustratively be obtained from updated bilingual dictionary 220, which was used during alignment.

[0070] FIG. 3C illustrates a target logical form 256 for the present example. It can be seen that the logical form segments from "click" to "button" and from "button" to "option" were stitched together from linked logical form 254 to obtain target logical form 256.

[0071] Generation component 228 is illustratively a rule-based, application-independent generation component that maps from target logical form 256 to the target string (or output target sentence) 258. Generation component 228 may illustratively have no information regarding the source language of the input logical forms, and works exclusively with information passed to it by transfer component 226. Generation component 228 also illustratively uses this information in conjunction with a monolingual (e.g., for the target language) dictionary to produce target sentence 258. One generic generation component 228 is thus sufficient for each language.

[0072] It can thus be seen that the present system parses information from various languages into a shared, common, logical form so that logical forms can be matched among different languages. The system can also utilize simple filtering techniques in building the transfer mapping database to handle noisy data input. Therefore, the present system can be automatically trained using a very large number of sentence pairs. In one illustrative embodiment, the number of sentence pairs is in excess of 10,000. In another illustrative embodiment, the number of sentence pairs is greater than 50,000 to 100,000, and may be in excess

of 180,000, 200,000, 350,000 or even in excess of 500,000 or 600,000 sentence pairs. Also, the number of sentence pairs can vary for different languages, and need not be limited to these numbers.

#### Logical Form Alignment

[0073] FIG. 4 illustrates a method 300 of associating logical forms of at least sentence fragments from two different languages, wherein the logical forms comprise nodes organized in a parent/child structure. Method 300 includes associating nodes of the logical forms to form tentative correspondences as indicated at block 302 and aligning nodes of the logical forms by eliminating at least one of the tentative correspondences and/or structural considerations as indicated at block 304.

[0074] As indicated above with respect to FIG. 2, alignment component 210 accesses bilingual dictionary 214 in order to form tentative correspondences, typically lexical correspondences, between the logical forms. Bilingual dictionary 214 can be created by merging data from multiple sources, and can also use inverted target-to-source dictionary entries to improve coverage. As used herein, bilingual dictionary 214 also represents any other type of resource that can provide correspondences between words. Bilingual dictionary 214 can also be augmented with translation correspondences acquired using statistical techniques.

[0075] In FIG. 2, the statistical techniques are performed by component 208. Although the output from component 208 can be used by alignment component 210, it is not necessary for operation of alignment component 210. However, one embodiment of component 208 will be described here, briefly, for the sake of completeness.

[0076] Component 208 receives a parallel, bilingual training corpus that is parsed into its content words. Word association scores for each pair of content words consisting of a word of language 1.1 that occurs in a sentence aligned in the bilingual corpus to a sentence of language 1.2 in which the other word occurs. A pair of words is considered "linked" in a pair of aligned sentences if one of the words is the most highly associated, of all the words in its sentence, with the other word. The occurrence of compounds is hypothesized in the training data by identifying maximal, connected sets of linked words in each pair of aligned sentences in the processed and scored training data. Whenever one of these maximal, connected sets contains more than one word in either or both of the languages, the subset of the words in that language is hypothesized as a compound. The original input text is rewritten, replacing the hypothesized compounds by single, fused tokens. The association scores are then recomputed for the compounds (which have been replaced by fused tokens) and any remaining individual words in the input text. The association scores are again recomputed, except that this time, co-occurrences are taken into account in computing the association scores only where there is no equally strong or stronger other association in a particular pair of aligned sentences in the training corpus.

[0077] Translation pairs can be identified as those word pairs or token pairs that have association scores above a threshold, after the final computation of association scores.

[0078] Similarly, component 208 can also assist in identifying translations of "captoids", by which we mean titles,

or other special phrases, all of whose words are capitalized. (Finding translations of capitals presents a special problem in languages like French or Spanish, in which convention dictates that only the first word of such an item is capitalized, so that the extent of the capital translation is difficult to determine.) In that embodiment, compounds are first identified in a source language (such as English). This can be done by finding strings of text where the first word begins with a capital letter, and later tokens in the contiguous string do not begin with a lowercase letter. Next, compounds are hypothesized in the target text by finding words that start with a capital letter and flagging this as the possible start of a corresponding compound. The target text is then scanned from left to right flagging subsequent words that are most strongly related to words in the identified compound in the source text, while allowing up to a predetermined number (e.g., 2) contiguous non-most highly related words, so long as they are followed by a most highly related word.

[0079] The left to right scan can be continued until more than the predetermined number (e.g., more than 2) contiguous words are found that are not most highly related to words in the identified compound in the source text, or until no more most highly related words are present in the target text, or until punctuation is reached.

[0080] While the above description has been provided for component 208, it is to be noted that component 208 is optional.

[0081] Referring again to method 300 in FIG. 4, generally, forming tentative correspondences in step 302 is aggressively pursued with the purpose of attempting to maximize the number of tentative correspondences formed between the logical forms. Accuracy of the tentative correspondences is not the most important criteria in step 302 because step 304 will further analyze the tentative correspondences and remove those that are determined to be incorrect.

[0082] Bilingual dictionary 214 represents direct translations used for forming tentative correspondences. However, in order to form additional tentative correspondences, derivational morphology can also be used. For example, translations of morphological bases and derivations, and base and derived forms of translations, can also be used to form tentative correspondences in step 302. Likewise, tentative correspondences can also be formed between nodes of the logical forms wherein one of the nodes comprises more lexical elements or words than the other node. For instance, as is common, one of the nodes can comprise a single word in one of the languages, while the other node comprises at least two words in the other language. Closely related languages such as English, Spanish, etc. also have word similarity (cognates) that can be used with fuzzy logic to ascertain associations. These associations can then be used to form tentative correspondences.

[0083] At this point, it may be helpful to consider an example of logical forms to be aligned. Referring to FIG. 5A, logical form 320 was generated for the sentence "En Información del hipervínculo, haga clic en la dirección del hipervínculo", while logical form 322 was generated for the English translation as "Under Hypertlink Information, click the Hypertlink address."

[0084] FIG. 5A further illustrates each of the tentative correspondences 323 identified in step 302. As an example

of the aggressive pursuit of tentative correspondences in step 302, in this example, each of the occurrences of "Hiper-vínculo" includes two different tentative correspondences with "Hypertlink Information" and "hypertlink" in the English logical form 322.

[0085] Referring now to step 304, the logical forms are aligned, which can include eliminating one or more of the tentative correspondences formed in step 302, and/or which can be done as a function of structural considerations of the logical forms. In one embodiment, step 304 includes aligning nodes of the logical forms as a function of a set of rules. In a further embodiment, each of the rules of the set of rules is applied to the logical forms in a selected order. In particular, the rules are ordered to create the most unambiguous alignments ("best alignments") first, and then, if necessary, to disambiguate subsequent node alignments. It is important to note that the order that the rules are applied in is not based upon the structure of the logical forms, i.e., top-down processing or bottom-up processing but rather, to begin with the most linguistically meaningful alignments, wherever they appear in the logical form. As such, this set of rules can be considered to be applied to the nodes of each of the logical forms non-linearly as opposed to linearly based upon the structure of the logical forms. Generally, the rules are intended to be language-neutral in order that they can be universally applied to any language.

[0086] FIG. 6 generally illustrates application of the set of rules to the logical forms as method 328. At step 330, each of the nodes of the logical forms is considered to be "unaligned" as opposed to "aligned". The set of rules is applied to the unaligned nodes irrespective of structure at step 332 to form aligned nodes. Therefore, it is desirable to distinguish between unaligned nodes and aligned nodes. One technique includes assigning all of the nodes initially to the set of unaligned nodes, and removing nodes when they are aligned. The use of sets whether actively formed in different locations of a computer readable medium or virtually formed through the use of Boolean tags associated with the nodes merely provides a convenient way in which to identify unaligned nodes and aligned nodes.

[0087] At step 332, the set of rules is applied to each of the unaligned nodes. FIG. 7 schematically illustrates aspects of step 332 that can be implemented to apply the set of rules. In one embodiment as discussed above, the rules are applied in a specified order. Herein "N" is a counter that is used to indicate which of the rules is applied. In the first iteration, step 334 applies the first rule to each of the unaligned nodes. If a rule fails to be applicable to any of the unaligned nodes, another rule from the set (and in one embodiment, the next successive rule indicative of a linguistically meaningful alignment) is then applied as indicated at steps 336 and 338.

[0088] If all the rules of the set of rules have been applied to all the nodes at step 340, the alignment procedure is finished. It should be noted that under some situations, not all of the nodes will be aligned.

[0089] If a rule can be applied to a set of nodes of the logical forms, the nodes are identified as being aligned and removed from the set of unaligned nodes, and application of the rules continues. However, in one embodiment, it is advantageous to begin again with the rules once some rules have been applied to obtain a more linguistically meaningful alignment. Therefore, it can be desirable to again apply rules

that have previously been applied. In this manner, in one embodiment, each of the rules of the set of rules is applied again starting with, for example, the first rule as indicated at step 342.

[0090] The following is an exemplary set of rules for aligning nodes of the logical forms. The set of rules presented herein is ordered based on the strongest to weakest linguistically meaningful alignments of the nodes. As appreciated by those skilled in the art, reordering of at least some of the rules presented herein may not significantly alter the quality of alignments of the logical forms.

[0091] 1. If a bi-directionally unique translation exists between a node or set of nodes in one logical form and a node or set of nodes in the other logical form, the two nodes or sets of nodes are aligned to each other. A bi-directionally unique translation exists if a node or a set of nodes of one logical form has a tentative correspondence with a node or a set of nodes in the other logical form, such that every node in the first set of nodes has a tentative correspondence with every node in the second set of nodes, and no other correspondences, and every node in the second set of nodes has a tentative correspondence with every node in the first set of nodes, and no other correspondences.

[0092] 2. A pair of parent nodes, one from each logical form, having a tentative correspondence to each other, are aligned with each other if each child node of each respective parent node is already aligned to a child of the other parent node.

[0093] 3. A pair of child nodes, one from each logical form, are aligned with each other if a tentative correspondence exists between them and if a parent node of each respective child node is already aligned to a corresponding parent node of the other child.

[0094] 4. A pair of nodes, one from each logical form, are aligned to each other if respective parent nodes of the nodes under consideration are aligned with each other and respective child nodes are also aligned with each other.

[0095] 5. A node that is a verb and an associated child node that is not a verb from one logical form are aligned to a second node that is a verb of the other logical form if the associated child node is already aligned with the second verb node, and either the second verb node has no aligned parent nodes, or the first verb node and the second verb node have child nodes aligned with each other.

[0096] 6. A pair of nodes, one from each logical form, comprising the same part-of-speech, are aligned to each other, if there are no unaligned sibling nodes, and respective parent nodes are aligned, and linguistic relationships between the set of nodes under consideration and their respective parent nodes are the same.

[0097] 7. A pair of nodes, one from each logical form, comprising the same part-of-speech, are aligned to each other if respective child nodes are aligned with each other and the linguistic relationship between the set of nodes under consideration and their respective child nodes are the same.

[0098] 8. If an unaligned node of one of the logical forms having immediate neighbor nodes comprising respective parent nodes, if any, all aligned, and respective child nodes, if any, all aligned, and if exactly one of the immediate nodes

is a non-compound word aligned to a node of the other logical form comprising a compound word, then align the unaligned node with the node comprising the compound word. Note that the immediate neighbor nodes herein comprise adjacent parent and child nodes however the existence of parent and child nodes is not required, but if they are present, they must be aligned.

[0099] 9. A pair of nodes, one from each logical form, comprising pronouns, are aligned to each other if respective parent nodes are aligned with each other and neither of the nodes under consideration have unaligned siblings.

[0100] 10. A pair of nodes, one from each logical form, comprising nouns are aligned to each other if respective parent nodes comprising nouns are aligned with each other and neither of the nodes under consideration have unaligned sibling nodes, and wherein a linguistic relationship between each of the nodes under consideration and their respective parent nodes comprises either a modifier relationship or a prepositional relationship.

[0101] 11. A first verb node of one logical form is aligned to a second verb node of the other logical form if the first verb node has no tentative correspondences and has a single associated child verb node that is already aligned with the second verb node.

[0102] 12. A first verb node and a single, respective parent node of one logical form is aligned to a second verb node of the other logical form if the first verb node has no tentative correspondences and has a single parent verb node that is already aligned with the second verb node, where the single parent verb node has no unaligned verb child nodes besides the first verb node, and the second verb node has no unaligned verb child nodes.

[0103] 13. A first node comprising a pronoun of one logical form is aligned to a second node of the other logical form if a parent node of the first node is aligned with the second node and the second node has no unaligned child nodes.

[0104] 14. A first verb node and a respective parent verb of one logical form is aligned to a second verb node of the other logical form if the first verb node has no tentative correspondences and the parent verb node is aligned with the second verb node and where the relationship between the first verb and the parent verb node comprise a modal relationship.

[0105] Some general classifications of the rules provided above include that one rule (rule 1) is primarily based upon the correspondences established in step 302, and in the embodiment illustrated, it is considered to be the strongest meaningful alignment since no ambiguity is present. Other rules such as rules 2, 3, 11, 12 and 14 are based on a combination of, or a lack of, tentative correspondences and the structure of the nodes under consideration and previously aligned nodes. The remaining rules rely solely on relationships between nodes under consideration and previously aligned nodes. Other general classifications that can be drawn include that the rules pertain to verbs, nouns and pronouns.

[0106] Referring back to the logical forms and tentative correspondences of FIG. 5A, the rules set out above can be applied according to the method 300 of FIG. 4 in order to



align the nodes as illustrated in FIG. 5B. In this example, the two instances of "Hipervinculo" have two ambiguous tentative correspondences, and while the correspondence from "Información" to "Hyperlink Information" is unique, the reverse is not. It should also be noted that neither the monolingual nor the bilingual lexicons or dictionaries have been customized for this domain. For example, there is no entry in the lexicon for "Hyperlink Information". This unit has been assembled by general rules that link sequences of capitalized words. Tentative lexical correspondences established for this element are based on translations found for its individual components.

[0107] Applying the alignment rules as described above, the alignment mappings created by the rules are illustrated in FIG. 5B as dotted lines 344, and are obtained as follows.

[0108] Iterating through the rules again, rule 1 applies in three places, creating alignment mappings between "dirección" and "address", "usted" and "you", and "click" and "click". These are the initial "best" alignments that provide the anchors from which the method will work outwards to align the rest of the structure.

[0109] Rule 2 does not apply to any nodes, but Rule 3 applies next to align the instance of "hipervinculo", that is the child of "dirección" to "hyperlink", which is the child of "address". The alignment method thus leveraged a previously created alignment ("dirección" to "address") and the structure of the logical form to resolve the ambiguity present at the lexical level.

[0110] Rule 1 applies (where previously it did not) to create a many-to-one mapping between "Información" and "hipervinculo" to "Hyperlink Information". The uniqueness condition in this rule is now met because the ambiguous alternative was cleared away by the prior application of Rule 3.

[0111] Rule 4 does not apply, but rule 5 applies to rollup "hacer" with its object "click", since the latter is already aligned to a verb. This produces the many-to-one alignment of "hacer" and "click" to "click".

[0112] Referring back to FIG. 7, alignment of the logical forms is completed when the rules are no longer applicable to any of the nodes. At this point, transfer mappings can be obtained by component 212.

[0113] FIG. 8 illustrates some of the transfer mappings obtainable from the example of aligned logical forms in FIG. 5B (other than transfer mapping 353 which is included as an example of a conflicting transfer mapping discussed in the next section). Generally, a transfer mapping or simply "mapping" is indicative of associating a word or logical form of a first language with a corresponding word or logical form of a second language. The mappings can be stored on any computer readable medium as explicit pointers linking the words or logical forms of the first language with the corresponding words or logical forms of the second language. Likewise, the mappings can be stored with the words or logical forms rather than in a separate database. As appreciated by those skilled in the art, other techniques can be used to associate words or logical forms of the first language with words or logical forms of the second language, and it is this association, that constitutes the mappings regardless of the specific techniques used in order to record this information.

[0114] Each mapping created during the alignment procedure can be a base structure upon which further mappings with additional context are also created. In particular, information can be stored on a computer readable medium to translate text from a first language to a second language, where the information comprises a plurality of mappings. Each mapping is indicative of associating a word or logical form of the first language with a word or logical form of the second language. However, in addition, at least some of the mappings corresponding to logical forms of the first language have varying context with some common elements. Likewise, at least some of the logical forms of the second language corresponding to the logical forms of the first language may also have varying context with some common elements. In other words, at least some of the core mappings obtained from the alignment procedure are used to create other, competing mappings having varying types and amounts of local context.

[0115] Referring to FIG. 8, mappings 350, 352, and 354 illustrate how an element of a logical form can vary. Mapping 350 comprises the base or core mapping on which further mappings are created. Mapping 352 expands the core mapping 350 to include an additional linguistic element, herein the direct object of the word "click", while the mapping 354 is expanded from the core mapping 350 such that the additional element comprises an under-specified node (" ") indicating a part of speech but no specific lemma. By comparing the mappings 350, 352 and 354, as well as mappings 356 and 358, it can be seen that the logical forms of the first language have common elements (parts of speech and/or lemmas), while the logical forms of the second language also have common elements.

[0116] By storing mappings indicative of logical forms with overlapping context, during translation run time, fluency and general applicability of the mappings for translating between the languages is maintained. In particular, by having mappings associating both words and smaller logical forms of the languages, translation from the first language to the second language is possible if the text to be translated was not seen in the training data. However, to the extent that the larger context was present in the training data, this is also reflected in the mappings such that when a mapping of larger context is applicable, a more fluent translation between the first language and the second language can be obtained.

[0117] Generally, linguistic constructs are used to provide boundaries for expanding the core mappings to include additional context. For example, a mapping for an adjective can be expanded to include the noun it modifies. Likewise, a mapping for a verb can be expanded to include the object as context. In another example, mappings for noun collocations are provided individually as well as a whole. As further illustrated in FIG. 8, some of the mappings can include under-specified nodes (" "), wherein the part of speech is indicated but no specific lemma is provided. These types of mappings increase the overall applicability of the mappings for translating from the first language to the second language, but also include context to enhance fluency of the translation obtained.

[0118] In general, mappings that can be created may have any number of wild-card or under-specified nodes, which may be under-specified in a number of different ways. For example, they may or may not specify a part-of-speech, and

they may specify certain syntactic or semantic features. For example, a pattern may have a wild-card node with the feature "ProperName" or "Location" marked, indicating that the pattern only applies when that node is matched to an input node that has the same feature. These wild-cards allow the system to hypothesize generalized mappings from specific data.

#### Matching the Transfer Mappings During Run Time

[0119] In addition to the information pertaining to the mappings between the words or logical forms of the first language and the second language, additional information can also be stored or used during run time translation. The additional information can be used to choose an appropriate set of mappings and resolve conflicts as to which mappings to use, i.e. (referring to FIG. 2) when a source logical form 252 (or part thereof) generated for a source sentence 250 matches more than one source side of the transfer mappings in the transfer mapping database 218.

[0120] For example, when the source logical form matches the source side of multiple transfer mappings in database 218, a subset of these matching transfer mappings is illustratively selected such that all transfer mappings in the subset are compatible with one another (i.e., they are not conflicting) and based on a metric that is a function of how much of the input sentence the transfer mappings in the subset collectively cover, as well as other measures related to individual transfer mappings. Some such measures are set out in Table 1.

TABLE 1

1.	Size of transfer mapping matched
2.	The frequency with which the transfer mapping was seen in the training data.
3.	The frequency with which the transfer mapping was generated from fully aligned logical forms.
4.	The frequency with which the transfer mapping was generated from partially aligned logical forms.
5.	The frequency with which the transfer mapping was generated from logical forms that resulted from a filled parse.
6.	An alignment score assigned to the transfer mapping by the alignment component.

[0121] Once the subset of matching transfer mappings is selected, the transfer mappings in the subset are combined into a transfer logical form from which the output text is generated.

[0122] It should be noted that the subset of matching transfer mappings can contain overlapping transfer mappings, so long as they are compatible. For example, the following logical form can be generated for the Spanish sentence "Haga clic en el direccion de la oficina" which can be translated as "Click the office address":

[0123] Hacer -- Dobj - click

[0124] - en - direccion

[0125] - de - oficina

[0126] This logical form can potentially be matched to all of the transfer mappings 350, 352 and 354 because each

transfer mapping contains this logical form. These transfer mappings overlap, but do not conflict (because all can be translated as the same thing). Therefore, all may be included in the subset of matching transfer mappings, and the transfer logical form can be generated from them. However, if it is desired to choose among them, the best choice may be transfer mapping 352 because it is the largest. Others could be chosen for a variety of other reasons as well.

[0127] An example of conflicting, matching transfer mappings is shown as transfer mapping 353, which conflicts with transfer mapping 352. Therefore, for example, the logical form:

[0128] Hacer -- Dobj - click

[0129] - en - direccion

[0130] would match all of transfer mappings 350, 352, 353 and 354. However, since transfer mappings 352 and 353 conflict (because they are translated differently) both cannot be part of the selected subset of matching transfer mappings. Thus, one is selected based on a predetermined metric. For example, subset 350, 352 and 354 can be compared against subset 350, 353 and 354 to see which covers the most nodes in the input logical form, collectively. Also, both transfer mappings 352 and 353 are the same size (on the source side). Therefore, other information can be used to distinguish between them in selecting the subset of matching transfer mappings.

[0131] As another example of conflicting transfer mappings, assume that a number of sentences processed during training included the phrase "click <something>" that aligned to the Spanish "hacer clic en <something>". In other sentences, assume the sentence "click <something>" aligned to "elegir <something>" (literally "select something").

[0132] This yields the following mappings (note these examples are English mapped to Spanish whereas previous examples have been Spanish mapped to English):

Click                      hacer  
Tobj -- \*                Tobj -- clic  
                                 en \*

[0133] for the first case, and

click                      elegir  
Tobj -- \*                Tobj -- \*  
                                 \*

[0134] in the second case.

[0135] In the proper context, translating "click" to "select" may be a legitimate variation. However it does present a problem in some cases. For example, notice that the source side of both transfers is identical, so at runtime, if the input logical form matches that source side, we are left with having to choose between the two different target sides, i.e. it must be decided whether to translate the input as "hacer clic ..." or as "elegir ...". In the absence of further context (which would likely have manifested itself by causing differing source sides of the transfers) we choose between them based on various frequency and scoring metrics.

[0136] Another type of conflict should also be mentioned. At runtime, for a given input sentence, there may be multiple matching transfer mappings that match different parts of the input sentence. Several of them can be chosen as the selected subset so that they can be stitched together to produce a transfer L.F. that covers the entire input. However, some of these matches that are stitched together will overlap one another, and some will not. Of the ones that overlap, we can only use those that are "compatible" with one another. As discussed above, by "overlap" we mean two mappings where at least one node of the input sentence is matched by both mappings. By compatible, we mean the following: matches are always compatible if they do not overlap, and matches that do overlap are compatible if the target sides that correspond to the node(s) at which they overlap are identical.

[0137] For example, if an input sentence is "cambiar configuraci3n de seguridad" (translated as "change the security setting") and it matches a transfer mapping as follows:

cambiar                      change  
Tobj -- configuraci3n      Tobj -- setting

[0138] and we match another mapping of:

configuraci3n              setting  
mod - seguridad          Mod security

[0139] then the two matches do overlap (on "configuraci3n"), but they are compatible, because they also both translate "configuraci3n" to "setting". Therefore, we can combine them to produce a transfer L.F. (or target L.F.):

[0140] change

[0141] Tobj setting

[0142] Mod security

[0143] However suppose there was also a third mapping of:

configuraci3n              value  
Mod - seguridad          Mod setting

[0144] then this mapping which does overlap the previous two at "configuraci3n", is not compatible, because it would translate "configuraci3n" to "value", not "setting". Therefore, this mapping cannot be merged with the previous two, so either this transfer mapping, or the previous two, must be chosen, but not both at the same time.

[0145] Table 1 shows examples of the information that can be used to further define the subset of matching transfer mappings (either to choose among conflicting matching transfer mappings or to narrow the subset of compatible, matching transfer mappings). Such information can include how much of the input sentence is covered by the subset of matching transfer mappings (collectively) and the size of the mappings, which can be ascertained from the logical form

that is matched in the transfer mapping itself. The size of a logical form includes both the number of specified nodes as well as the number of linguistic relationships between the nodes. Thus, by way of example, the size of the logical form from the source side of mapping 350 equals 2, while the size of the logical form on the target side equals 1. In another example, the logical form on the source side of mapping 354 equals 4, while the target side of mapping 354 equals 2.

[0146] The information for choosing the subset of transfer mappings can also include other information related to individual transfer mappings, such as the frequency with which the logical forms in the transfer mapping are seen in the training data. If desired, the training data can include "trusted" training data, which can be considered more reliable than other training data. The frequency of the mapping as seen in the trusted training data can be retained in addition, or in the alternative, to storing the frequency of the mapping as seen in all of the training data.

[0147] Other information that can be helpful in selecting the subset of matching transfer mappings includes the extent of complete alignment of the logical forms in the training data from which the logical forms of a transfer mapping have been obtained. In other words, the alignment procedure can fully or completely align the nodes of the larger logical forms, or some nodes can remain unaligned. In the example of FIG. 5H, all the nodes were aligned; however, as indicated above, this may not always be the case. Those mappings associated with fully aligned logical forms may be considered more reliable. Of course, information for resolving conflicts or further defining the subset can also indicate the frequency with which the mapping was generated from both fully aligned logical forms as well as partially aligned logical forms.

[0148] Likewise, additional information can include the frequency with which the logical forms in the transfer mapping originated from a complete parse of the corresponding training data. In particular, the frequency with which the mapping originated from a complete or fitted parse, or in contrast, the frequency that the mapping originated from only a partial parse can be stored for later use in resolving conflicts while matching during translation.

[0149] Another form of information can include a score or value assigned to the transfer mapping by the alignment procedure used to extract the mapping. For instance, the score can be a function of how "strong" (linguistically meaningful) the aligned nodes are (or how confident the alignment component is in the transfer mapping). The score can therefore be a function of when (which iteration) and which rule formed the alignment. The particular function or metric used to calculate the alignment score is not crucial, and any such metric can be used to generate information related to an alignment score that can be used during run time translation.

[0150] It should be noted that, although the present invention is described above primarily with respect to analyzing, aligning and using logical forms, at least some of the inventive concepts discussed herein are applicable to other dependency structures as well.

[0151] Although the present invention has been described with reference to particular embodiments, workers skilled in the art will recognize that changes may be made in form and detail without departing from the spirit and scope of the invention.

What is claimed is:

1. A computer implemented method of translating a textual input in a first language to a textual output in a second language, comprising:

generating an input logical form based on the textual input;

selecting a set of one or more of a plurality of matching transfer mappings in a transfer mapping database that match at least a portion of the input logical form, based on a predetermined metric;

combining the set of transfer mappings into a target logical form; and

generating the textual output based on the target logical form.

2. The method of claim 1 wherein the input logical form includes a plurality of input nodes and wherein selecting comprises:

selecting the set of transfer mappings based on a number of input nodes covered by the set of transfer mappings, collectively.

3. The method of claim 1 wherein selecting comprises:

selecting the set of transfer mappings based on sizes of the plurality of matching transfer mappings.

4. The method of claim 3 wherein selecting comprises:

selecting the set of transfer mappings as a largest of the plurality of matching transfer mappings.

5. The method of claim 1 wherein selecting comprises:

selecting the set of transfer mappings based on frequencies with which the plurality of matching transfer mappings were generated during a training phase used in training the transfer mapping database.

6. The method of claim 1 wherein selecting comprises:

selecting the set of transfer mappings based on frequencies with which the plurality of matching transfer mappings were generated from completely aligned logical forms during a training phase used in training the transfer mapping database.

7. The method of claim 1 wherein selecting comprises:

selecting the set of transfer mappings based on frequencies with which the plurality of matching transfer mappings were generated from partially aligned logical forms during a training phase used in training the transfer mapping database.

8. The method of claim 1 wherein selecting comprises:

selecting the set of transfer mappings based on frequencies with which the plurality of matching transfer mappings were generated from non-fitted parses of training data used to generate logical forms during a training phase used in training the transfer mapping database.

9. The method of claim 1 wherein selecting comprises:

selecting the set of transfer mappings based on a score associated with each of the plurality of matching transfer mappings, the score being indicative of a confidence in the transfer mapping with which it is associated.

10. The method of claim 1 wherein combining the set of transfer mappings comprises:

generating a linked logical form, indicative of links between the input logical form and logical forms in the transfer mapping database, based on the set of transfer mappings.

11. The method of claim 10 wherein combining further comprises:

generating a target logical form based on the linked logical form.

12. The method of claim 11 wherein generating the target logical form comprises:

accessing a bilingual dictionary based on words in the linked logical form.

13. The method of claim 11 wherein generating the textual output comprises:

generating the textual output based on the target logical form.

14. The method of claim 1 wherein selecting comprises:

selecting as the set a plurality of overlapping, matching transfer mappings.

15. The method of claim 14 wherein combining comprises:

combining the plurality of overlapping, matching transfer mappings to obtain the target logical form.

16. A machine translation system for translating a textual input in a first language to a textual output in a second language, the machine translation system comprising:

a matching component configured to match input logical forms generated based on the textual input to a set of one or more of a plurality of matching transfer mappings in a transfer mapping database that match at least a portion of the input logical form, based on a predetermined metric; and

a generation component configured to generate the textual output based on the selected transfer mappings.

17. The machine translation system of claim 16 and further comprising:

a logical form generator receiving the textual input and generating the input logical forms based on the textual input.

18. A machine translation system for translating a textual input in a first language to a textual output in a second language, the machine translation system comprising:

an input generator generating an input dependency structure based on the textual input;

a transfer mapping database including a plurality of transfer mapping dependency structures formed based on at least ten thousand parallel, aligned, training sentences;

a matching component configured to receive the input dependency structure and match it against a matching set of one or more of the transfer mapping dependency structures in the transfer mapping database; and

a generation component configured to generate the textual output based on the matching transfer mapping dependency structure.

19. The machine translation system of claim 18 wherein the transfer mapping database includes a plurality of transfer mapping dependency structures formed based on at least fifty thousand parallel, aligned, training sentences.

20. The machine translation system of claim 19 wherein the transfer mapping database includes a plurality of transfer mapping dependency structures formed based on at least one hundred thousand parallel, aligned, training sentences.

21. The machine translation system of claim 19 wherein the transfer mapping database includes a plurality of transfer mapping dependency structures formed based on at least one hundred eighty thousand parallel, aligned, training sentences.

22. The machine translation system of claim 19 wherein the transfer mapping database includes a plurality of transfer mapping dependency structures formed based on at least two hundred thousand parallel, aligned, training sentences.

23. A computer implemented method of training a transfer mapping database, comprising:

receiving a plurality of parallel, aligned, pairs of input sentences in two different languages;

generating input logical forms for the input sentences in both languages, the input logical forms being shared across both languages; and

training the transfer mapping database based on the input logical forms.

24. The method of claim 23 wherein training comprises: aligning the input logical forms to obtain transfer mappings; and

training the transfer mapping database based on the transfer mappings.

25. The method of claim 24 wherein training the transfer mapping database based on transfer mappings comprises:

training the transfer mapping database based only on transfer mappings obtained from the aligned logical forms at least a predetermined threshold number of times.

26. The method of claim 25 wherein the predetermined threshold number of times comprises two times.

27. The method of claim 23 wherein receiving comprises receiving at least ten thousand parallel, aligned, training sentences.

28. The method of claim 23 wherein receiving comprises receiving at least fifty thousand parallel, aligned, training sentences.

29. The method of claim 23 wherein receiving comprises receiving at least one hundred thousand parallel, aligned, training sentences.

30. The method of claim 23 wherein receiving comprises receiving at least one hundred eighty thousand parallel, aligned, training sentences.

31. The method of claim 23 wherein receiving comprises receiving at least two hundred thousand parallel, aligned, training sentences.

32. A computer implemented method of training a transfer mapping database, comprising:

receiving a plurality of parallel, aligned, pairs of input sentences in two different languages;

generating input logical forms for the input sentences in both languages;

aligning the input logical forms to obtain transfer mappings;

filtering the transfer mappings obtained; and

training the transfer mapping database based only on the filtered transfer mappings.

33. The method of claim 32 wherein filtering the transfer mappings comprises:

filtering transfer mappings obtained from the aligned logical forms less than at least a predetermined threshold number of times.

34. The method of claim 33 wherein filtering the transfer mappings comprises:

filtering transfer mappings obtained from the aligned logical forms less than at least two times.

\* \* \* \* \*

# **EXHIBIT B**

U.S. Patent No. 5,477,451

filed July 5, 2001

Brown et al.



US005477451A

**United States Patent** [19]

Brown et al.

[11] Patent Number: 5,477,451

[45] Date of Patent: Dec. 19, 1995

[54] **METHOD AND SYSTEM FOR NATURAL LANGUAGE TRANSLATION**

[75] Inventors: Peter F. Brown, New York; John Cocke, Bedford; Stephen A. Della Pietra, Pearl River; Vincent J. Della Pietra, Blauvelt; Frederick Jelinek, Briarcliff Manor; Jennifer C. Lai, Garrison; Robert L. Mercer, Yorktown Heights, all of N.Y.

[73] Assignee: International Business Machines Corp., Yorktown Heights, N.Y.

[21] Appl. No.: 736,278

[22] Filed: Jul. 25, 1991

[51] Int. Cl.<sup>6</sup> ..... G06F 17/20; G06F 17/27

[52] U.S. Cl. .... 364/419.08

[58] Field of Search ..... 364/419, 419.02, 364/419.08, 419.16; 381/43, 51

[56] **References Cited****U.S. PATENT DOCUMENTS**

4,754,489	6/1988	Bokser .....	382/40
4,829,580	5/1989	Church .....	381/52
4,852,173	7/1989	Bahl et al. ....	381/43
4,882,759	11/1989	Bahl et al. ....	381/51
4,984,178	1/1991	Hemphill et al. ....	364/419
4,991,094	2/1991	Fagen et al. ....	364/419
5,033,087	7/1991	Bohl et al. ....	381/43
5,068,789	11/1991	Van Vliembergen .....	364/419
5,109,509	4/1992	Katayama et al. ....	395/600
5,146,405	9/1992	Church .....	364/419
5,200,893	4/1993	Ozawa et al. ....	364/419

**FOREIGN PATENT DOCUMENTS**

0327266	8/1989	European Pat. Off. .
0357344	3/1990	European Pat. Off. .
0399533	11/1990	European Pat. Off. .
WO9010911	9/1990	WIPO .

**OTHER PUBLICATIONS**

"Method for Inferring Lexical Associations from Textual Co-Occurrences", IBM Technical Disclosure Bulletin, vol. 33 1B, Jun. 1990.

"Tagging Text with a Probabilistic Model", by Bernard Meriadlo, *Proceedings of the International Conference on Acoustics, Speech and Signal Processing*, Paris, France, May 14-17, 1991.

(List continued on next page.)

Primary Examiner—David M. Huntley

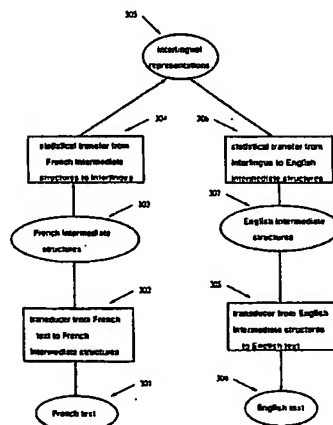
Assistant Examiner—Charles R. Kyle

Attorney, Agent, or Firm—Sterne, Kessler, Goldstein & Fox; Robert P. Tassinari

[57] **ABSTRACT**

The present invention is a system for translating text from a first source language into a second target language. The system assigns probabilities or scores to various target-language translations and then displays or makes otherwise available the highest scoring translations. The source text is first transduced into one or more intermediate structural representations. From these intermediate source structures a set of intermediate target-structure hypotheses is generated. These hypotheses are scored by two different models: a language model which assigns a probability or score to an intermediate target structure, and a translation model which assigns a probability or score to the event that an intermediate target structure is translated into an intermediate source structure. Scores from the translation model and language model are combined into a combined score for each intermediate target-structure hypothesis. Finally, a set of target-text hypotheses is produced by transducing the highest scoring target-structure hypotheses into portions of text in the target language. The system can either run in batch mode, in which case it translates source-language text into a target language without human assistance, or it can function as an aid to a human translator. When functioning as an aid to a human translator, the human may simply select from the various translation hypotheses provided by the system, or he may optionally provide hints or constraints on how to perform one or more of the stages of source transduction, hypothesis generation and target transduction.

42 Claims, 54 Drawing Sheets



## OTHER PUBLICATIONS

"Word-Sense Disambiguation Using Statistical Methods", by Peter F. Brown et al., appearing in the *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics*, Jun. 1991, pp. 264-270.

"Lex—A Lexical Analyzer Generator", M. E. Lesk, *Computer Science Technical Report*, No. 39, Bell Laboratories, Oct. 1975.

"Self Organized Language Modeling for Speech Recognition" by F. Jelinek, *Language Processing for Speech Recognition*, pp. 450-506.

"A Tree-Based Statistical Language Model for Natural Language Speech Recognition" by Lalit R. Bahl et al., *IEEE Transactions on Acoustics*, vol. 37, No. 7, Jul. 1989, pp. 1001-1008.

"Trainable Grammars for Speech Recognition" by James K. Baker, *Speech Communications Papers Presented at the 97th Meeting of the Acoustic Society of America*, 1979, pp. 547-550.

"Interpolated Estimation of Markov Source Parameters from

Sparse Data" by F. Jelinek and R. L. Mercer, *Workshop on Pattern Recognition in Practice*, Amsterdam (Netherlands): North Holland, May 21-23, 1980.

"An Inequality and Associated Maximization Technique in Statistical Estimation for Probabilistic Functions of Markov Processes", by Leonard E. Baum, *Inequalities*, vol. 3, 1972, pp. 1-8.

"Aligning Sentences in Parallel Corpora" by Peter F. Brown et al., appearing in the *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics*, Jun. 1991, pp. 169-176.

"Partial Traceback in Continuous Speech Recognition" by James C. Spohrer et al., *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, Paris, France, 1982.

"Deriving Translation Data from Bilingual Texts", Catizone et al., *Proceedings of the First International Acquisition Workshop*, Detroit, Mich., 1989.

"Making Connections", by M. Kay, appearing in *ACH/ALLC '91*, Tempe, Ariz., 1991, p. 1.



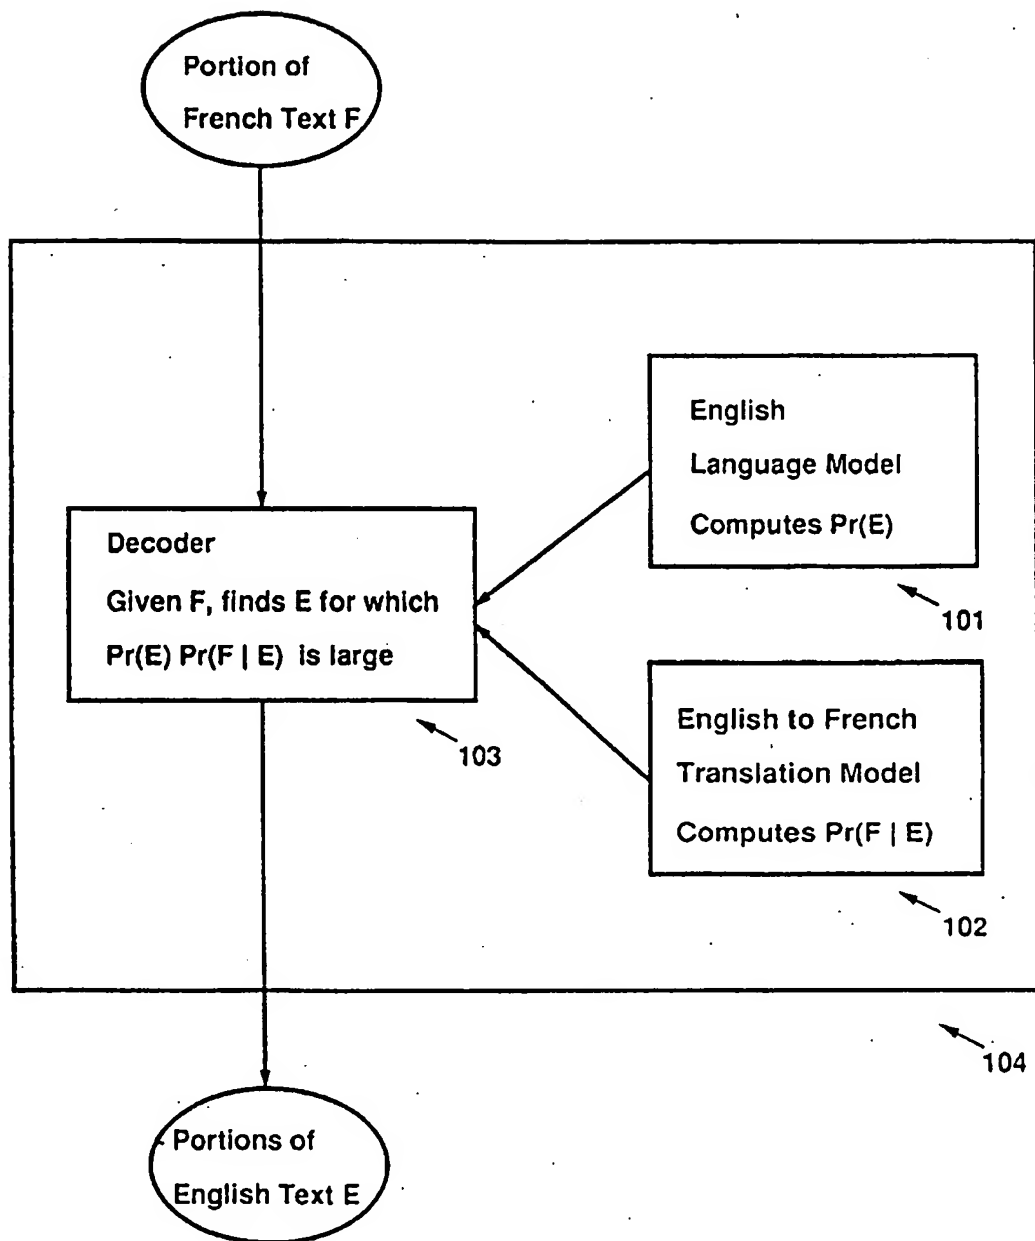


Figure 1

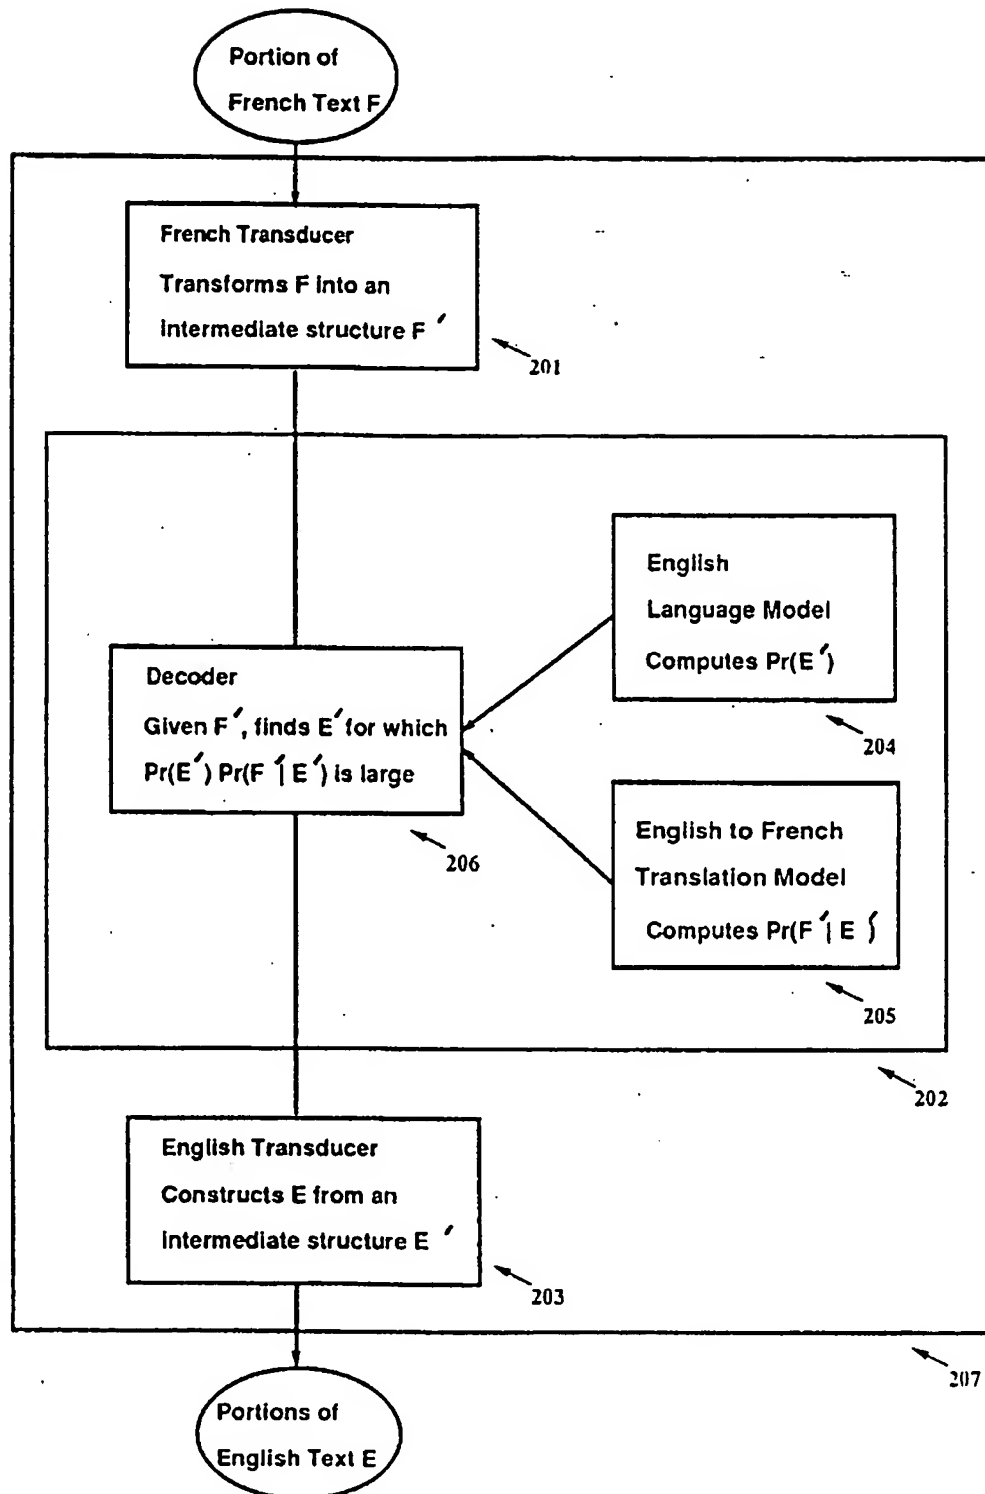


Figure 2

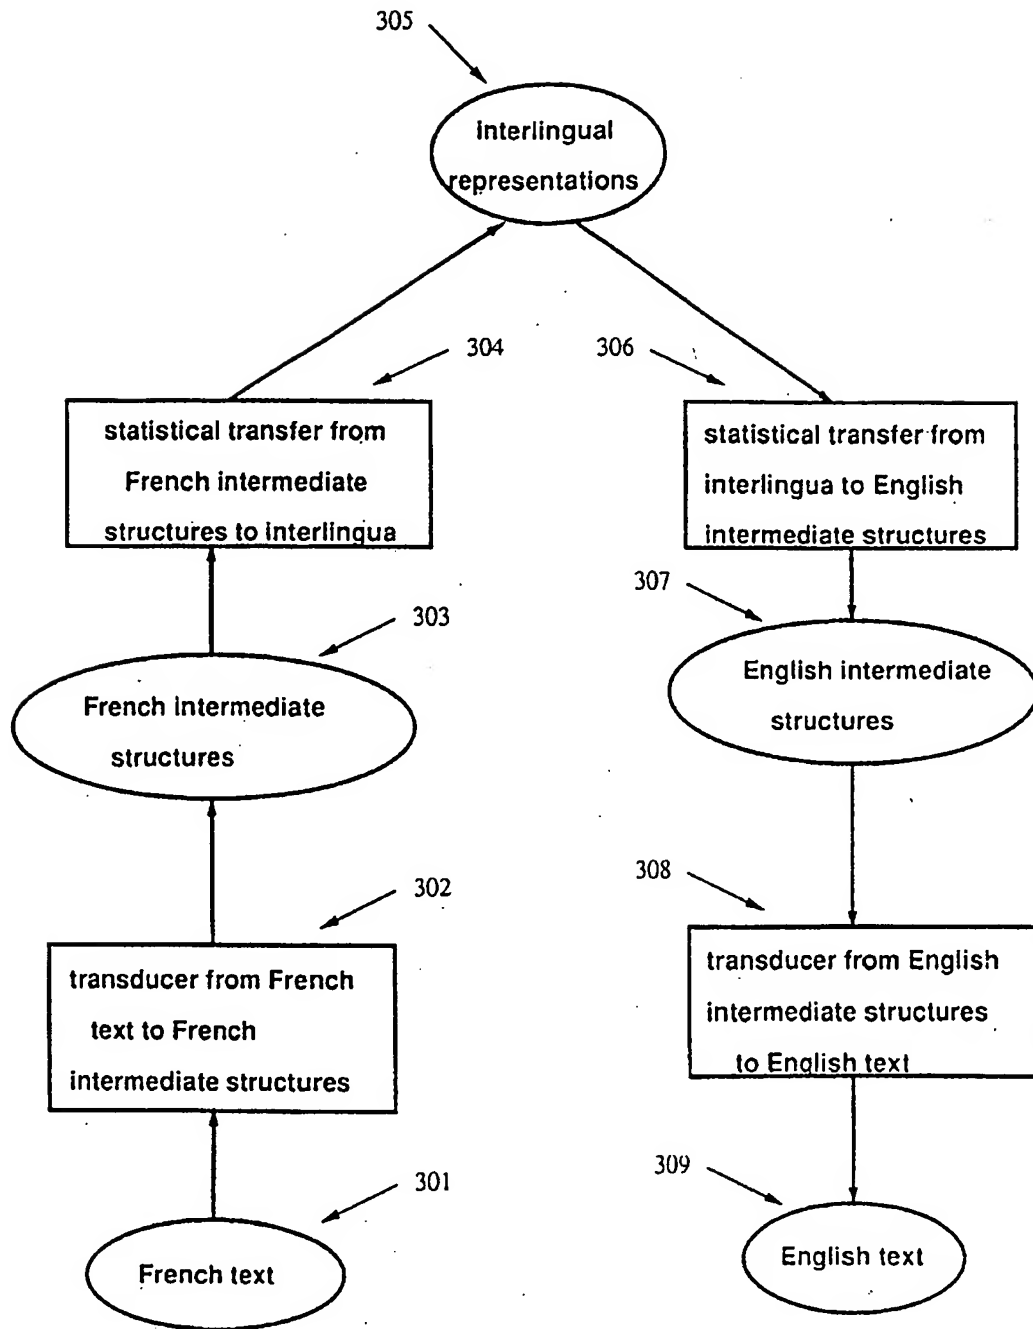


Figure 3

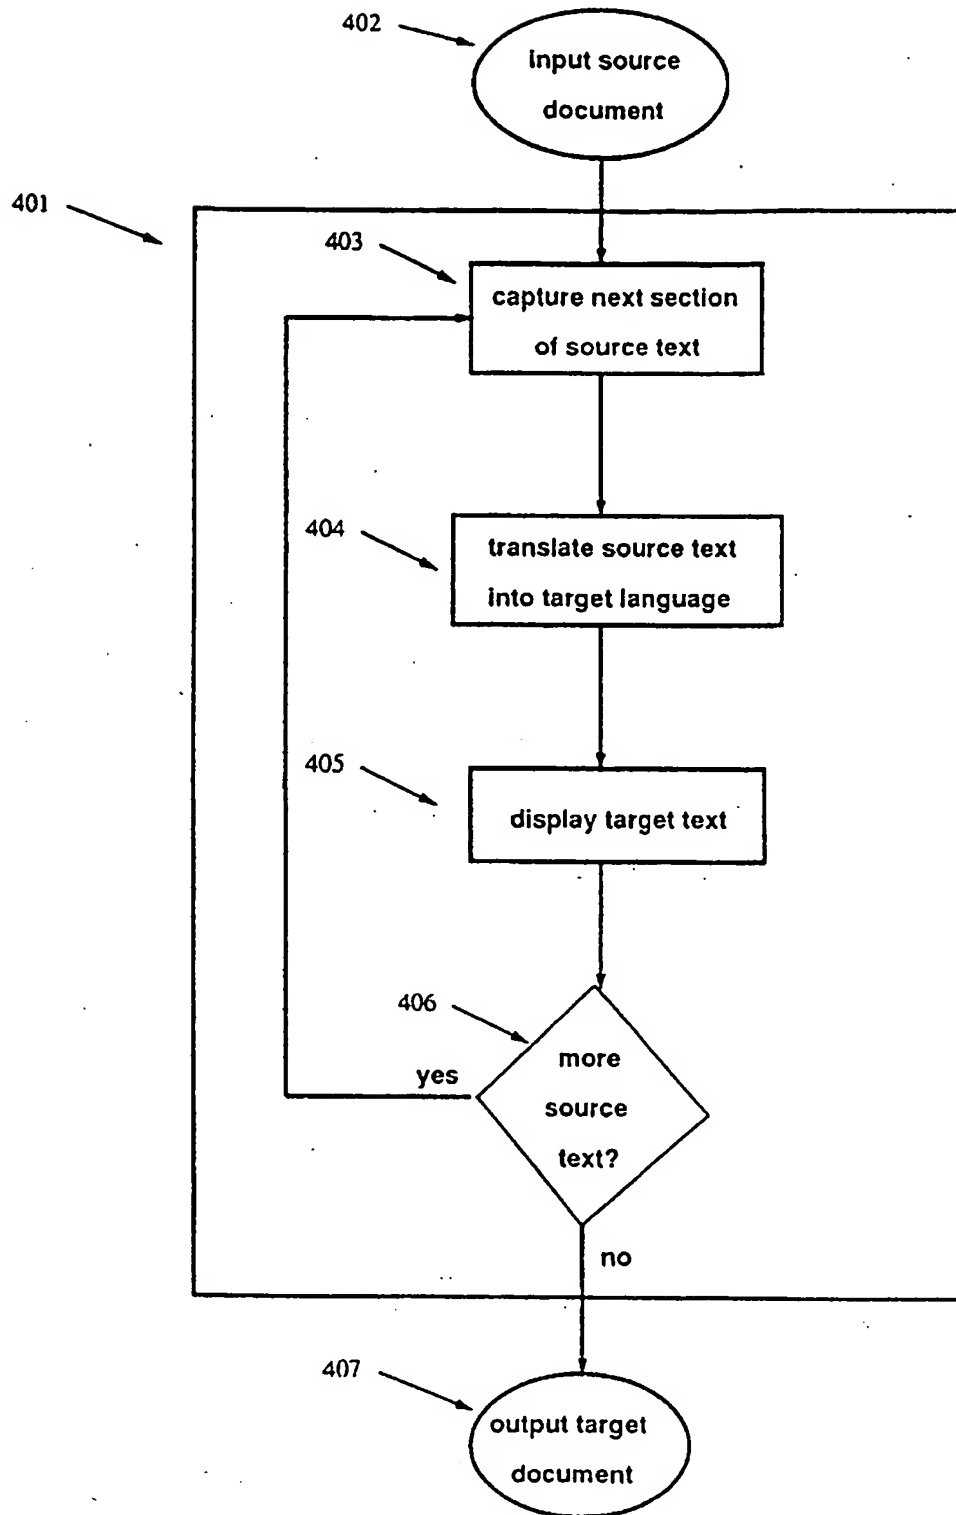


Figure 4

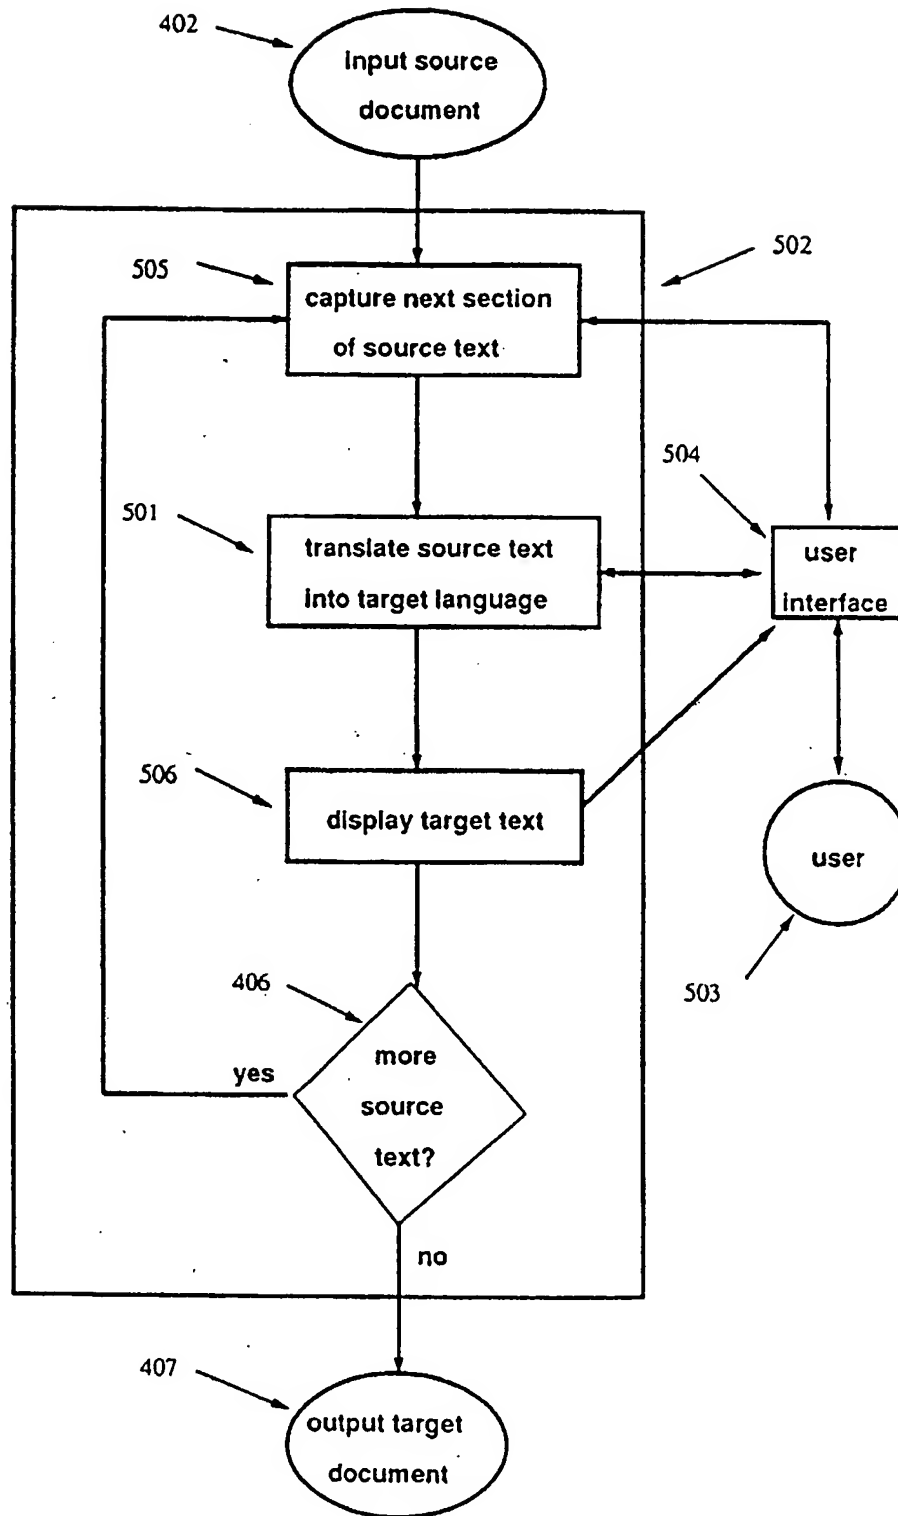


Figure 5




Figure 6

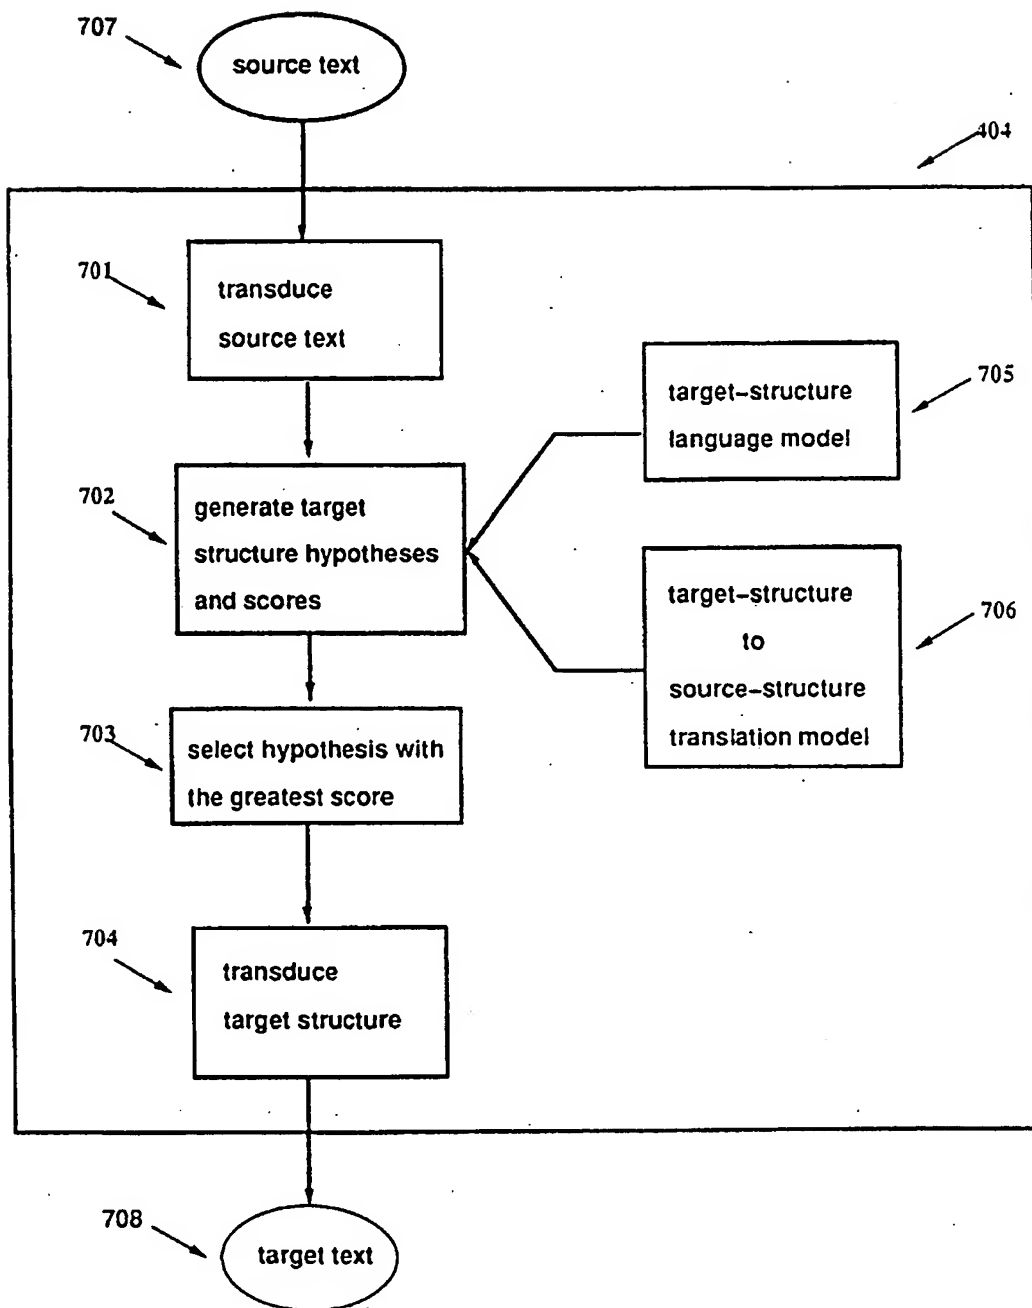


Figure 7

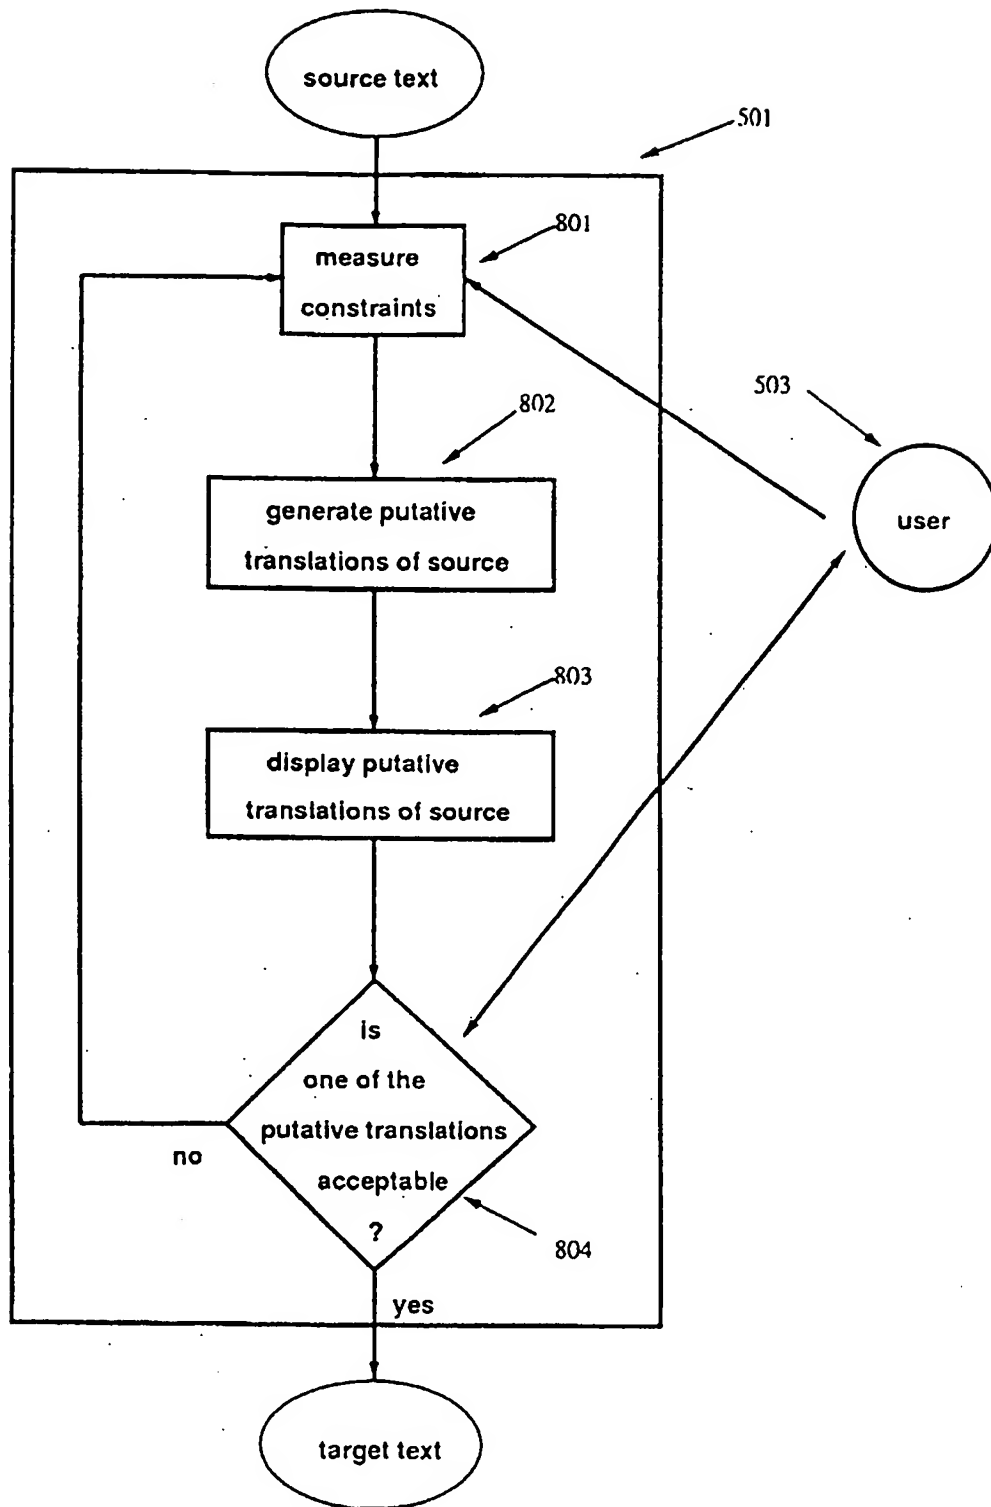


Figure 8



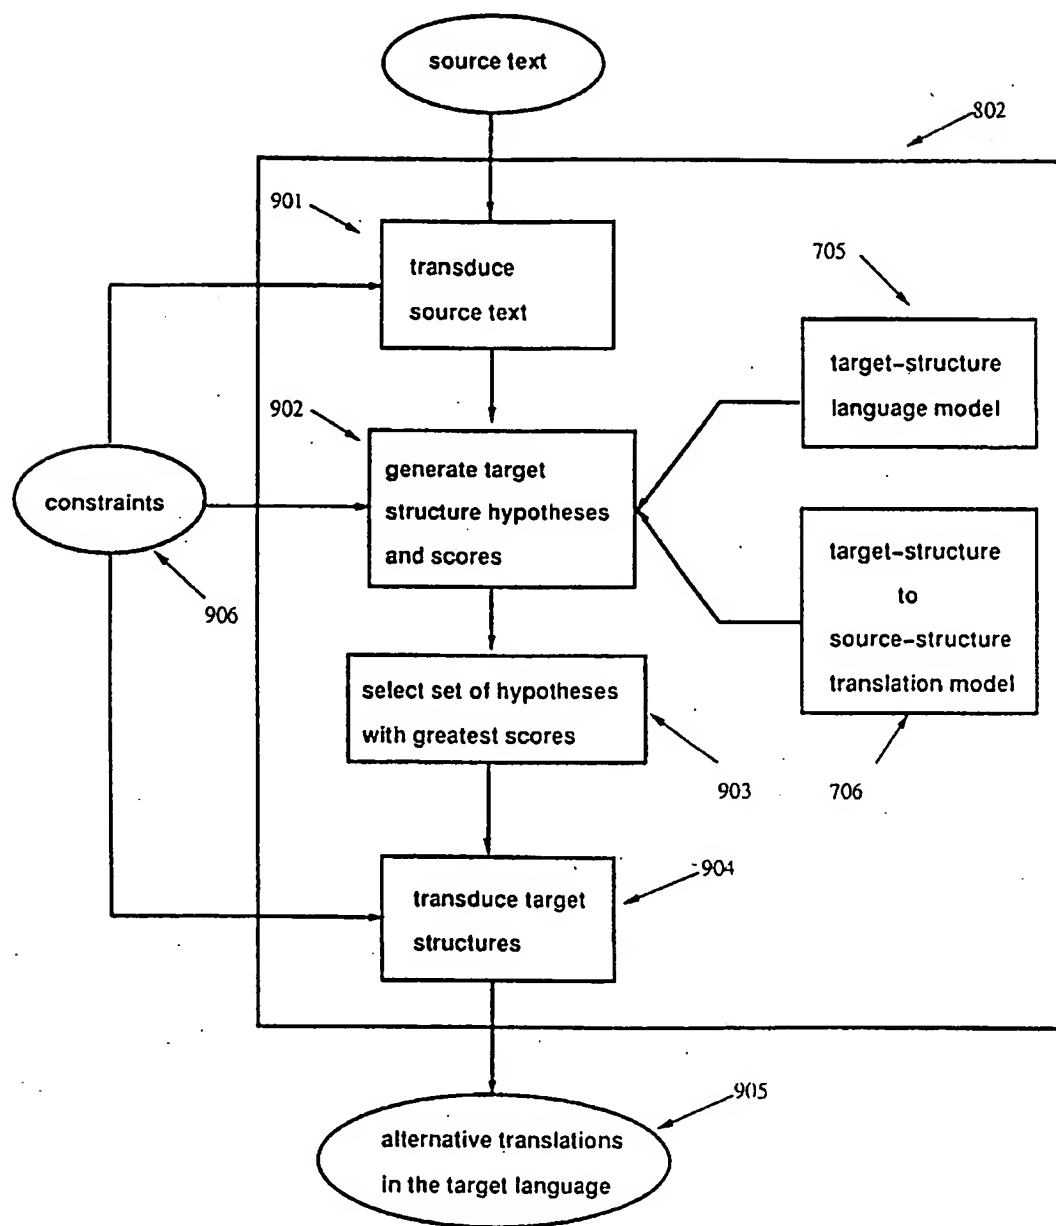


Figure 9




Figure 10

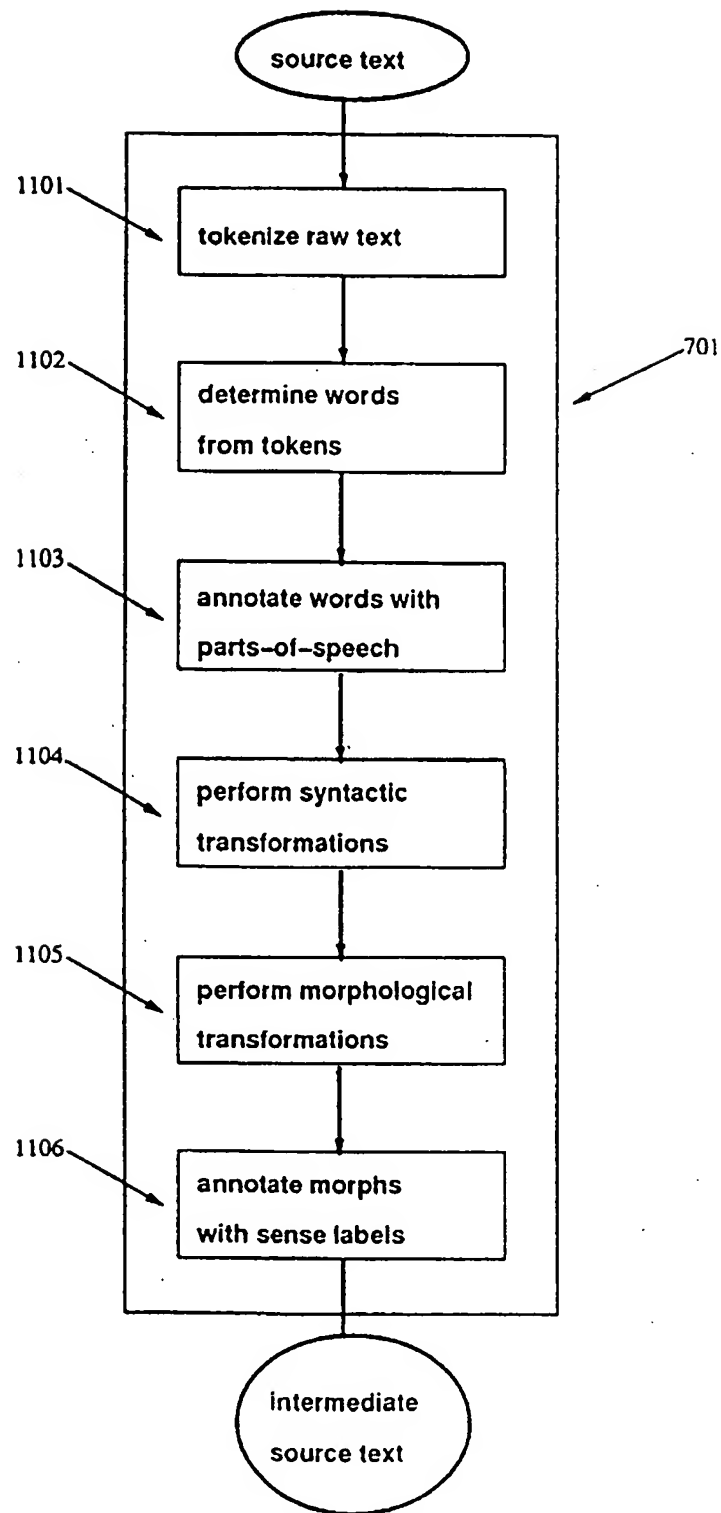


Figure 11

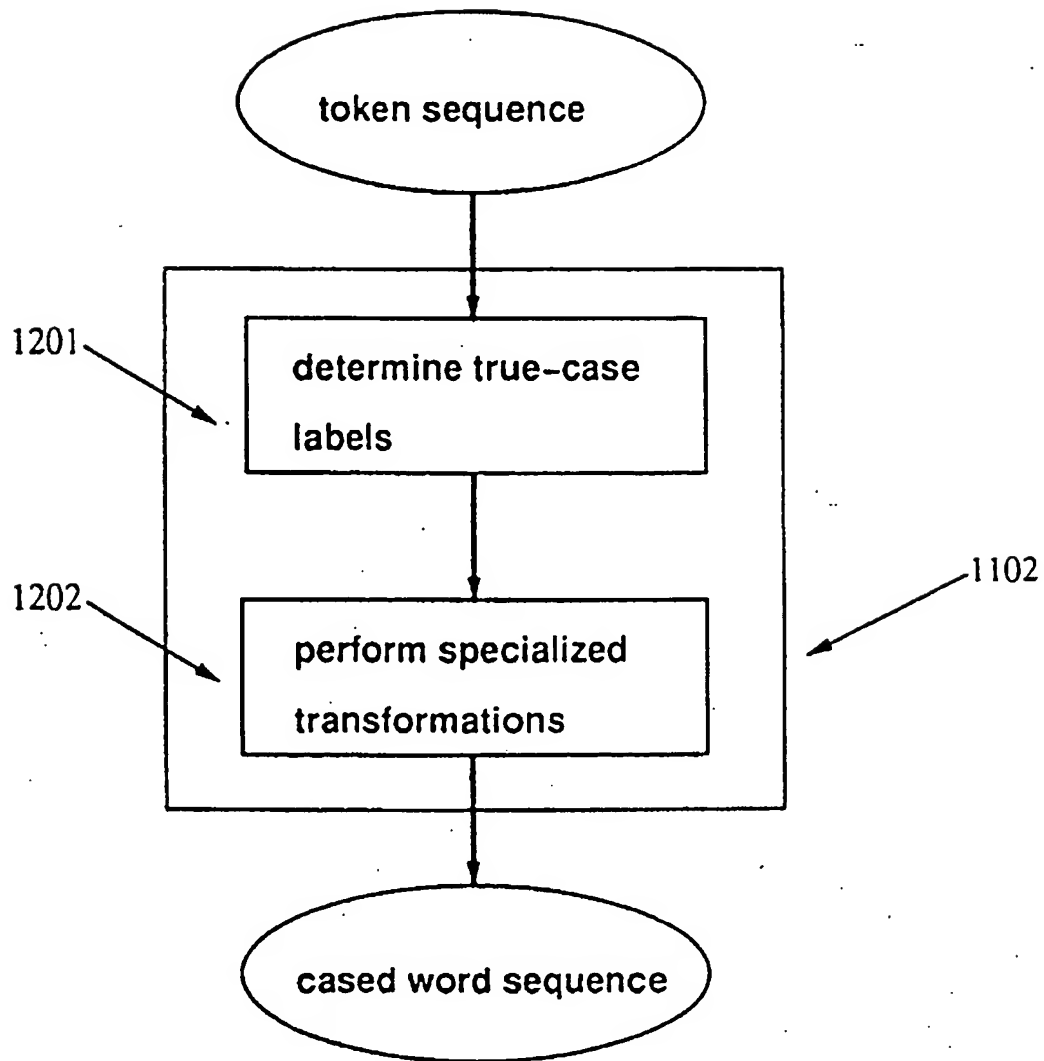


Figure 12




Figure 13




Figure 14

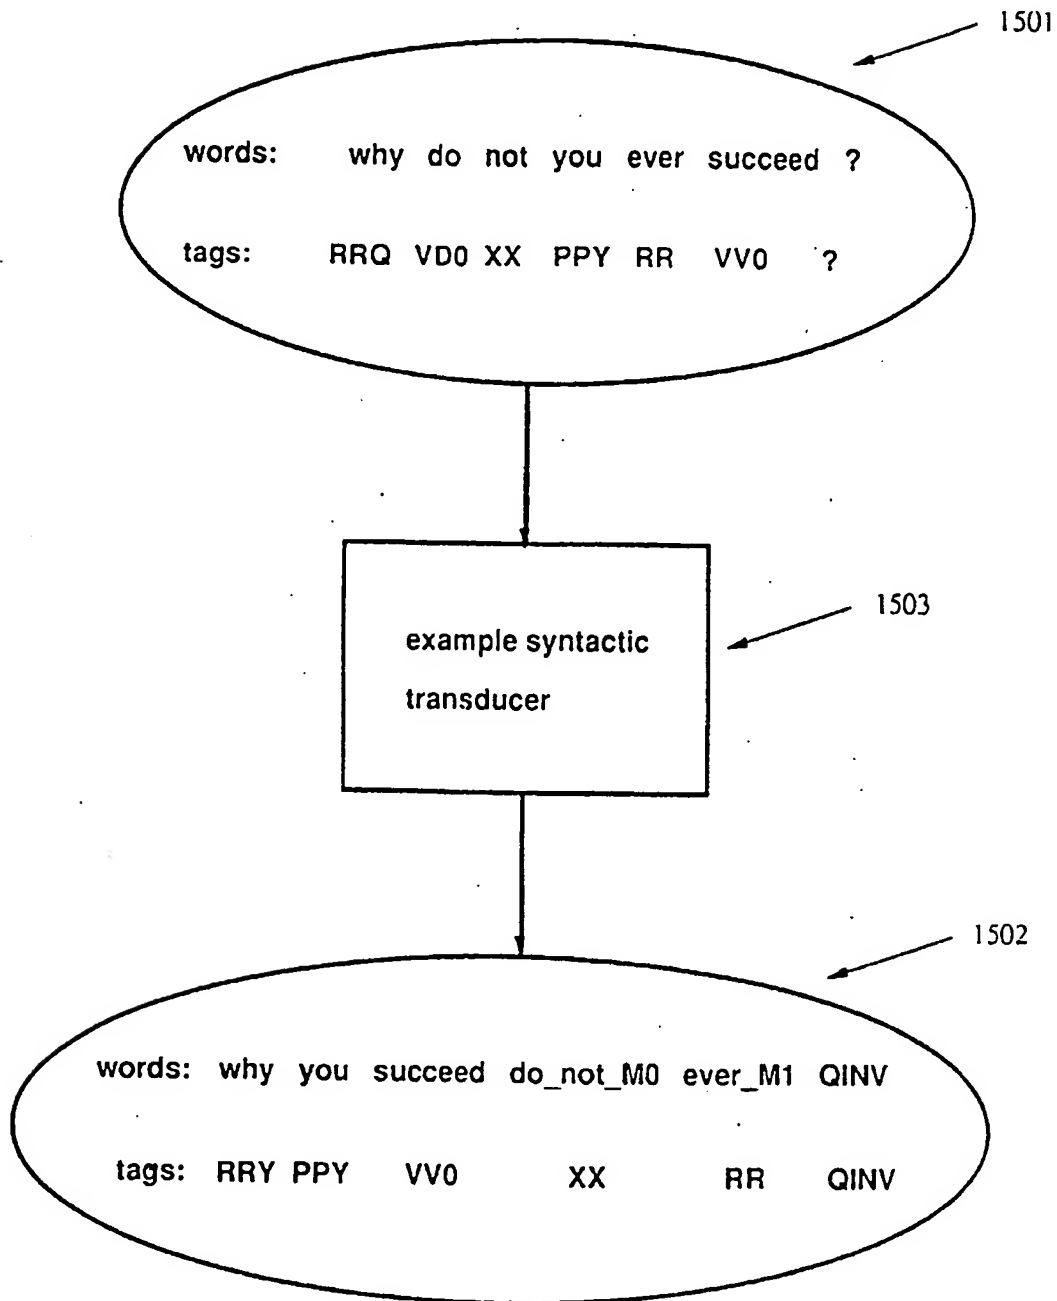


Figure 15

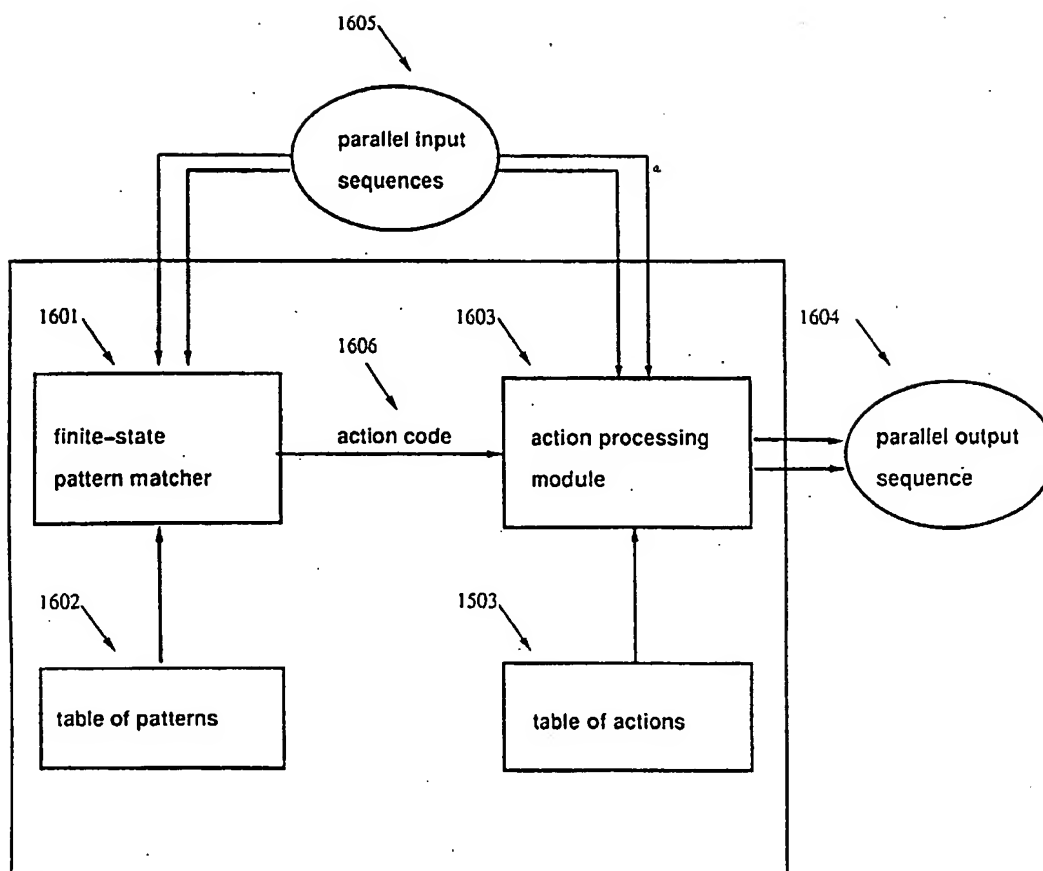


Figure 16



```

begin patterns and actions;
WH_NP?.[1].DO_SET.'not'?.[2].SUBJECT_NP.[3].ADVERB*.[4].
,BARE_INF_TAG.(#)*.'?'
-> english_question_inversion;
(#) +
-> default_action;
end patterns and actions;

begin auxiliary patterns;
ADVERB      = ,ADV_TAG | 'most','RR!' | 'more','RR!';
BARE_NP     = ((ADVERB)*.,ADJ_TAG)*.,COMMON_NOUN_TAG;
SUBJECT_NP  = (,SUBJECT_TAG) | (,DETERMINER_TAG)?.BARE_NP;
WH_NP      = WH_WORD | (( 'which' | 'what' | 'how'.'many').BARE_NP);
end auxiliary patterns;

begin sets of words;
DO_SET      = {do does did};
WH_WORD     = {what who where why whom how when which};
end sets of words;

begin sets of tags;
PROPER_NOUN_TAG = {NP NP1 NP2};
PRONOUN_TAG    = {PN PN1 PNQS};
SUBJECT_TAG    = {PROPER_NOUN_TAG PRONOUN_TAG};
BARE_INF_TAG   = {VVO VHO VBO VDO};
COMMON_NOUN_TAG = {ND1 NN NN1 NN2 NNJ NNJ1 NNJ2>NNL>NNL1>NNL2
NO NNO2 NNS1 NNS2 NNSA1 NNSB NNSB1
NNSB2 NNT1 NNT2 NNU NNU1 NNU2 NPD1 NPD2 NPM1};
ADV_TAG       = {RA REX RG RGA RGR RGT RP RPK RR RRR RRT RT};
ADJ_TAG       = {JA JB JBR JBT JJ JJR JJT JK};
DETERMINER_TAG = {AT AT1 DA DA1 DA2 DA2R DA2T DAR DAT DB DB2 DD
DD1 DD2 DDQ DDQ$ DDQV};
end sets of tags;

```

Figure 17

```
1  english_question_inversion: procedure;
2    /* This procedure is invoked when the pattern:
3       WH_NP?.[1].DO_SET.'not'?.[2].SUBJECT_NP.[3].ADVERB*. [4]
4       .#,BARE_INF_TAG.(#)*.'?' is matched */
5
6    set the output sequence to null;
7
8    append tuples the beginning of the input sequence
9      to position @reg(1)-1 to output;
10
11   append tuples from input positions
12     @reg(2) to @reg(3)-1 to output;
13
14   if @reg(2) = @reg(1) + 2 then do;
15     append tuples from input positions @reg(1) to
16       @reg(1)+1 to output;
17     append tuples from input in positions
18       @reg(3) to @reg(4) to output;
19   end;
20
21   else do;
22     append tuples from input in positions
23       @reg(3) to @reg(4)-1 to output;
24
25     if the word at postion @reg(1) is 'does' then do;
26       append the input word at position @reg(4) and the
27         tag for 'does' to output;
28     end;
29
30     else if the word at postion @reg(1) is 'do' then do;
31       append tuple at input position @reg(4)
32     end;
33
34     else if the word at postion @reg(1) is 'did' then do;
35       append the input word at position @reg(4)
36         and the tag for 'did' to output;
37     end;
38
39     else there is an error;
40
41   end;
42
43   append the tuples from input positions @reg(4)+1 to
44     one less than the last position to output;
45
46   append the word 'QINV' and the tag 'QINV' to output;
47
48   end english_question_inversion;
```

Figure 18

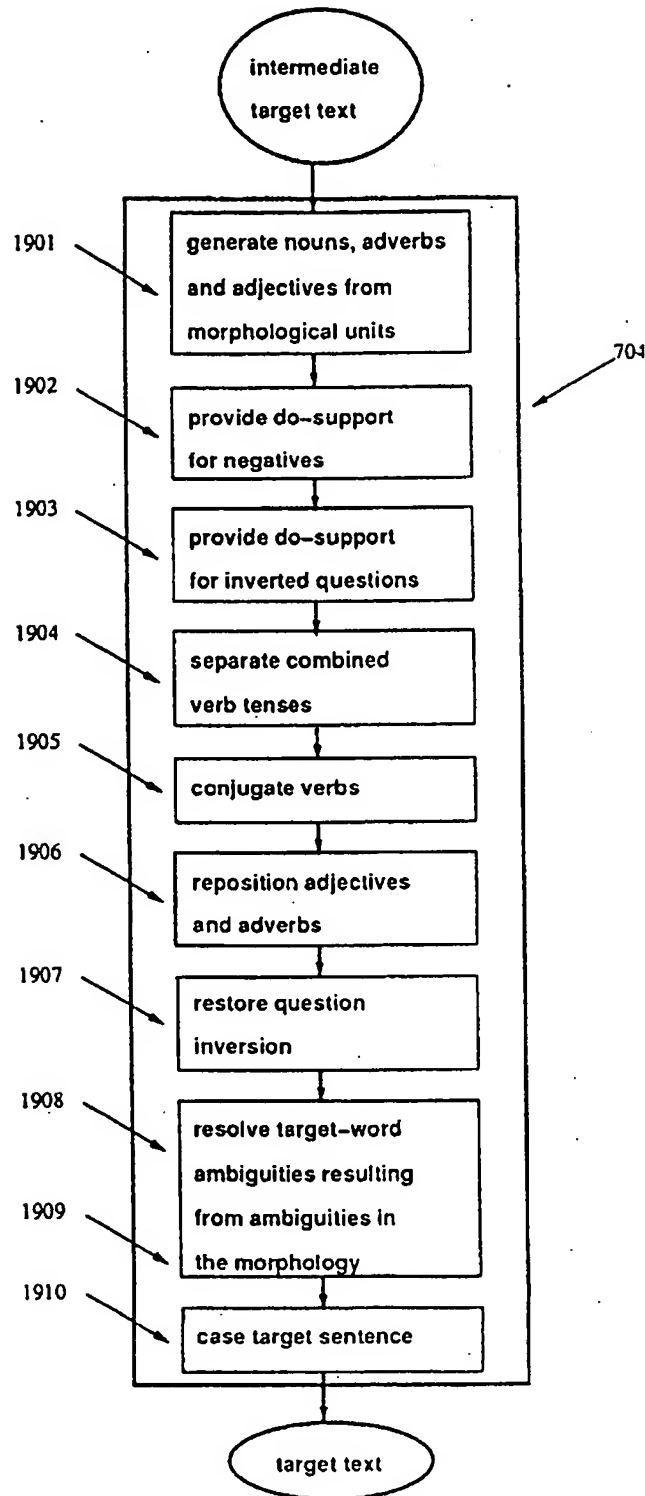


Figure 19




Figure 20

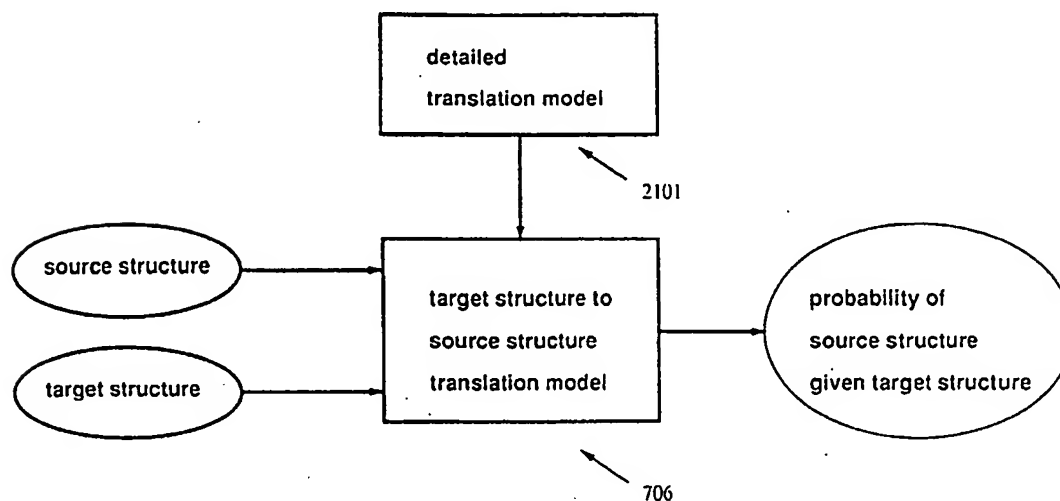


Figure 21




Figure 22




Figure 23

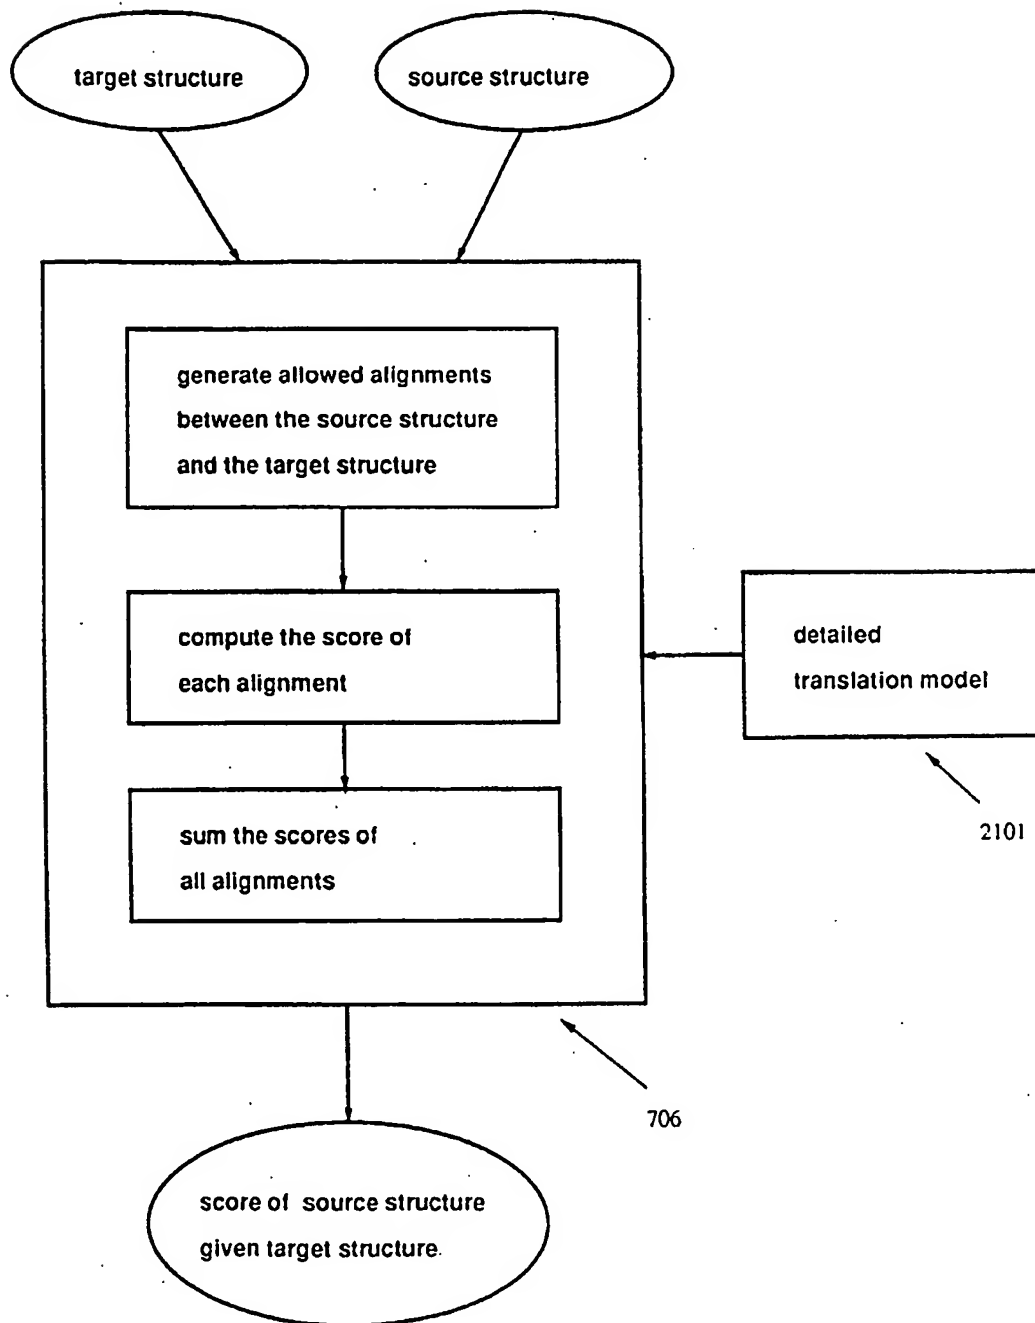


Figure 24

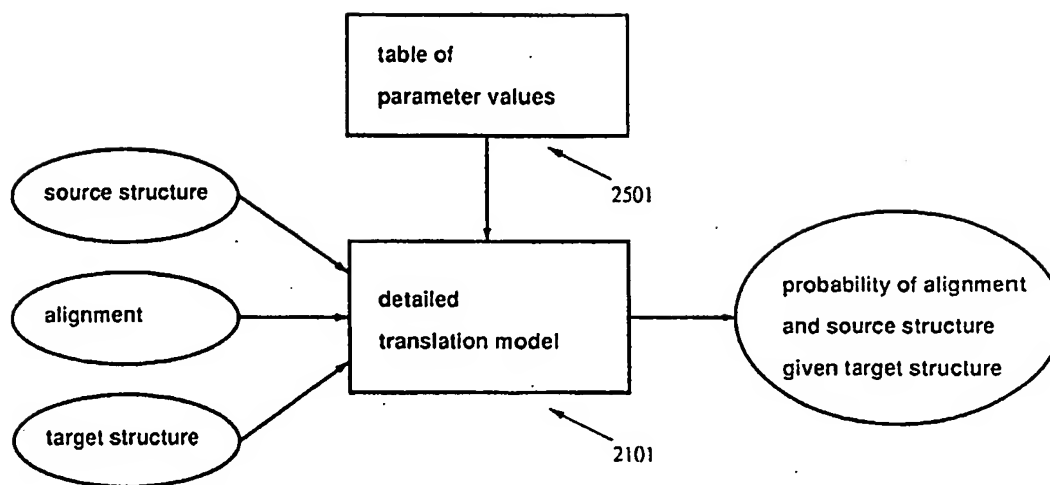


Figure 25

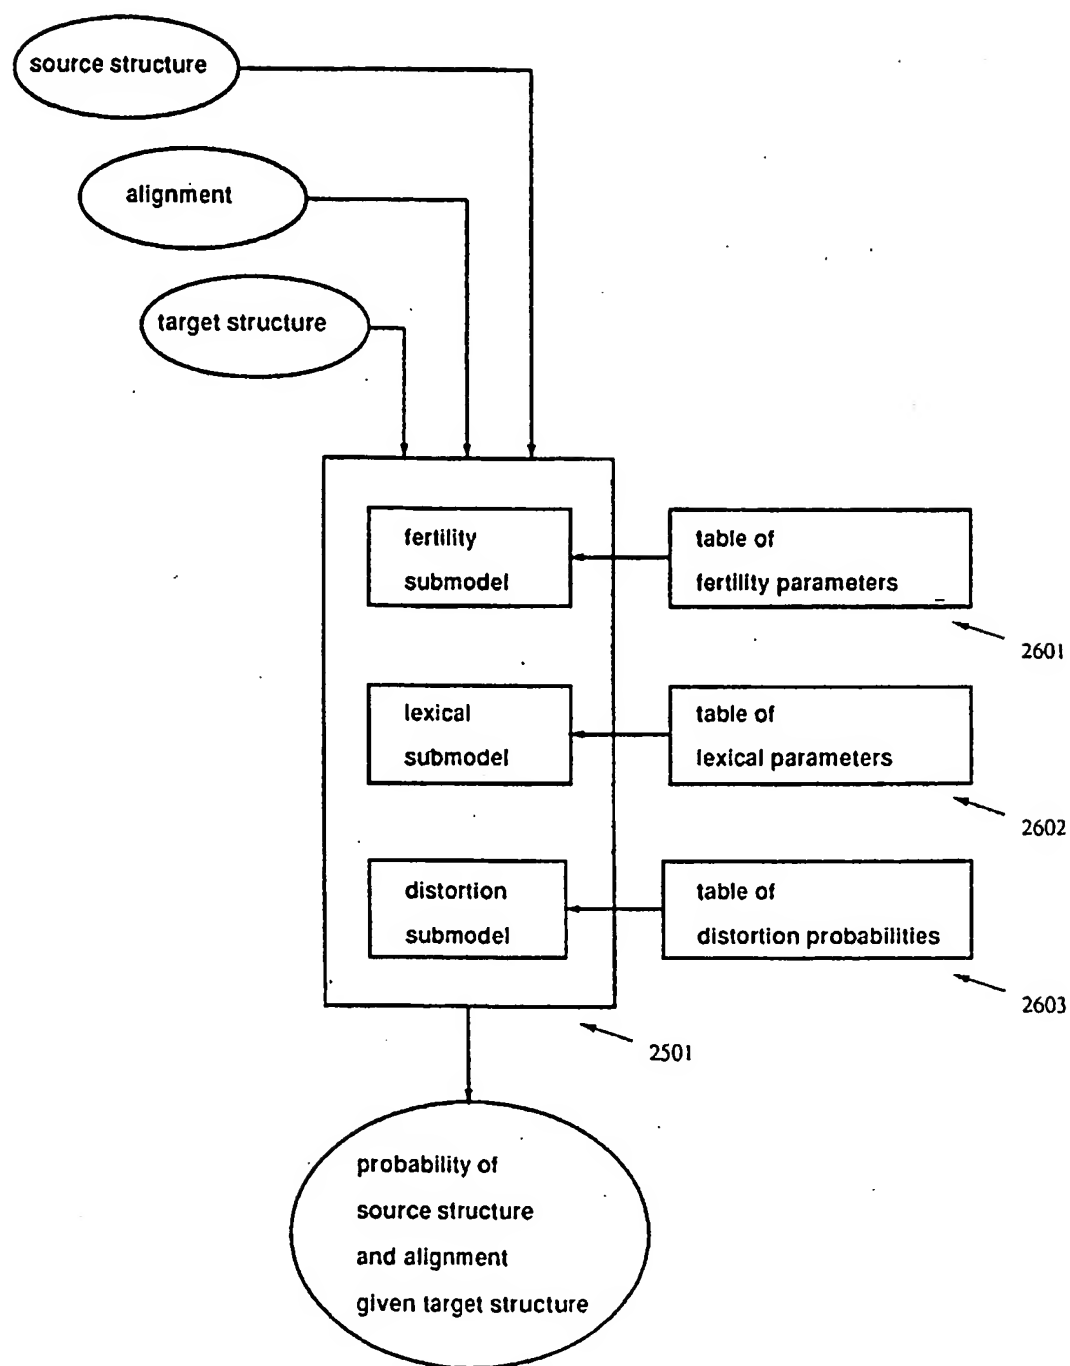


Figure 26






Figure 27

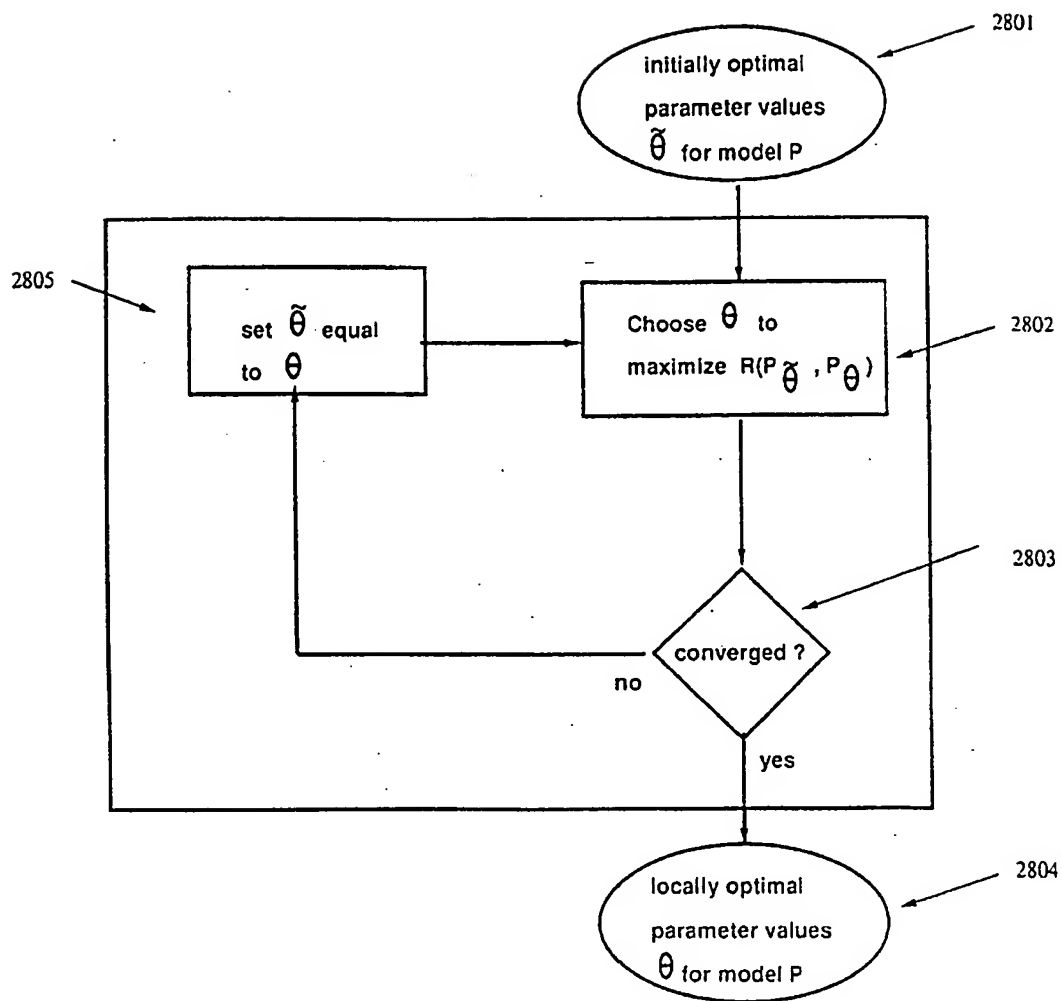


Figure 28

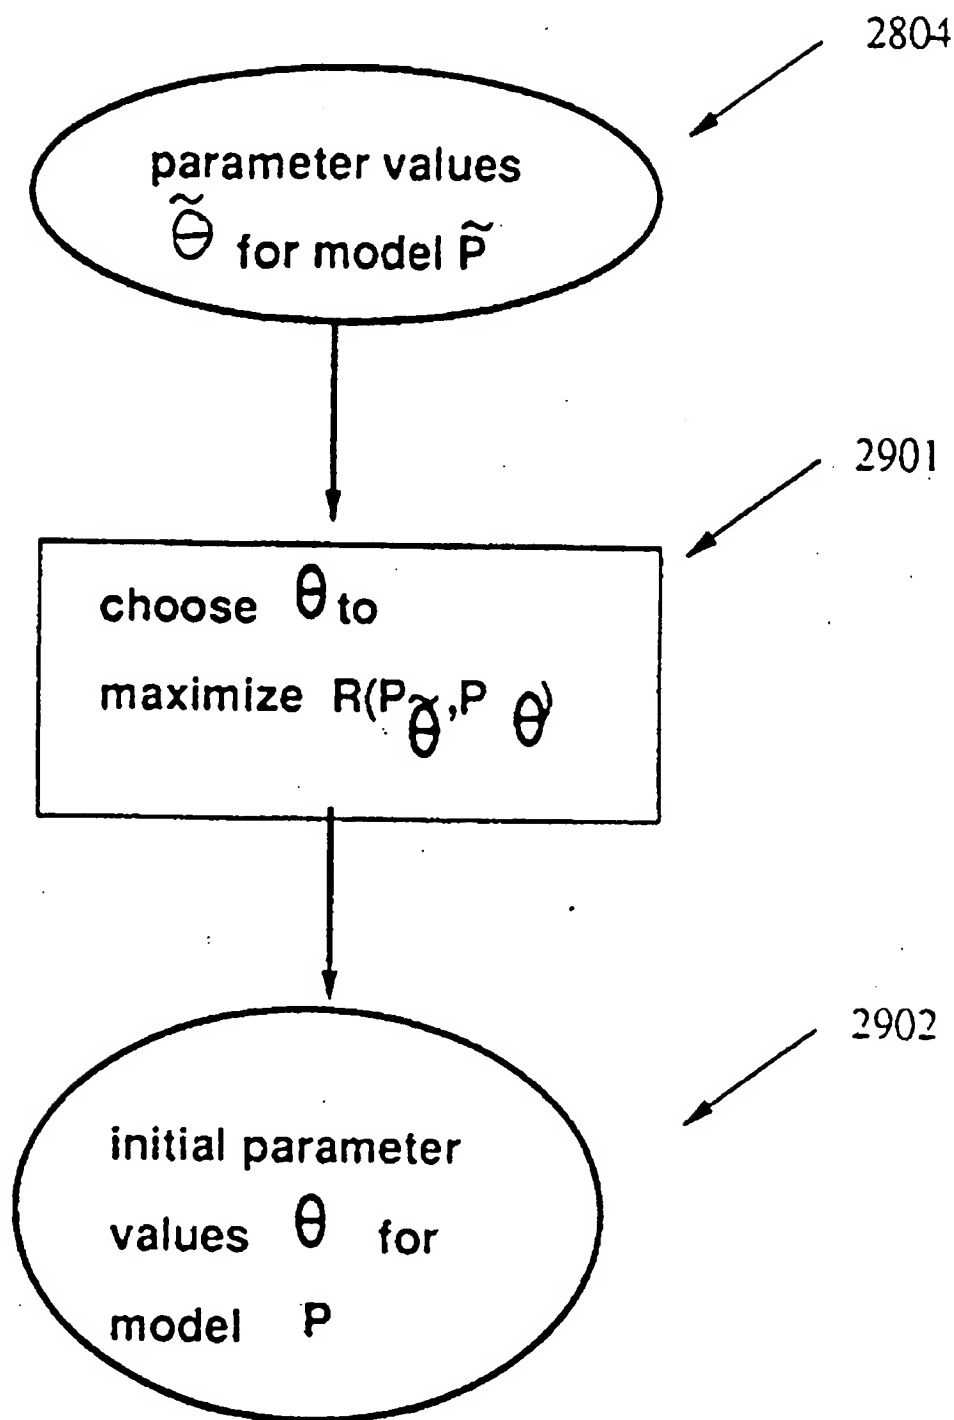


Figure 29

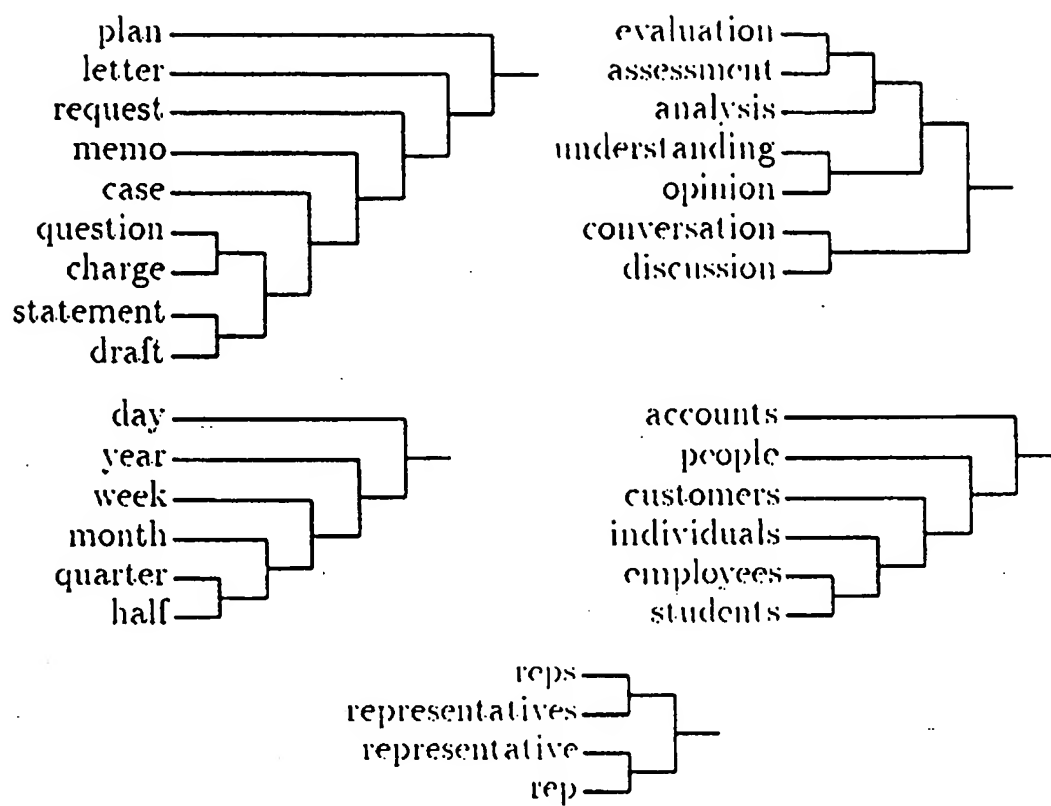


Figure 30

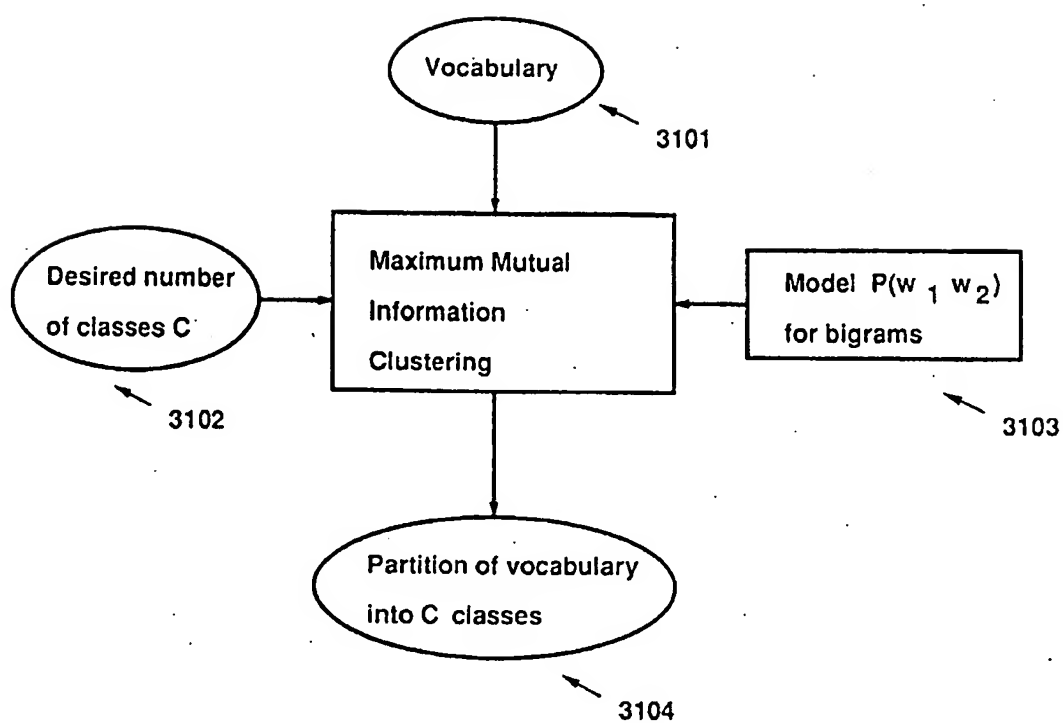


Figure 31

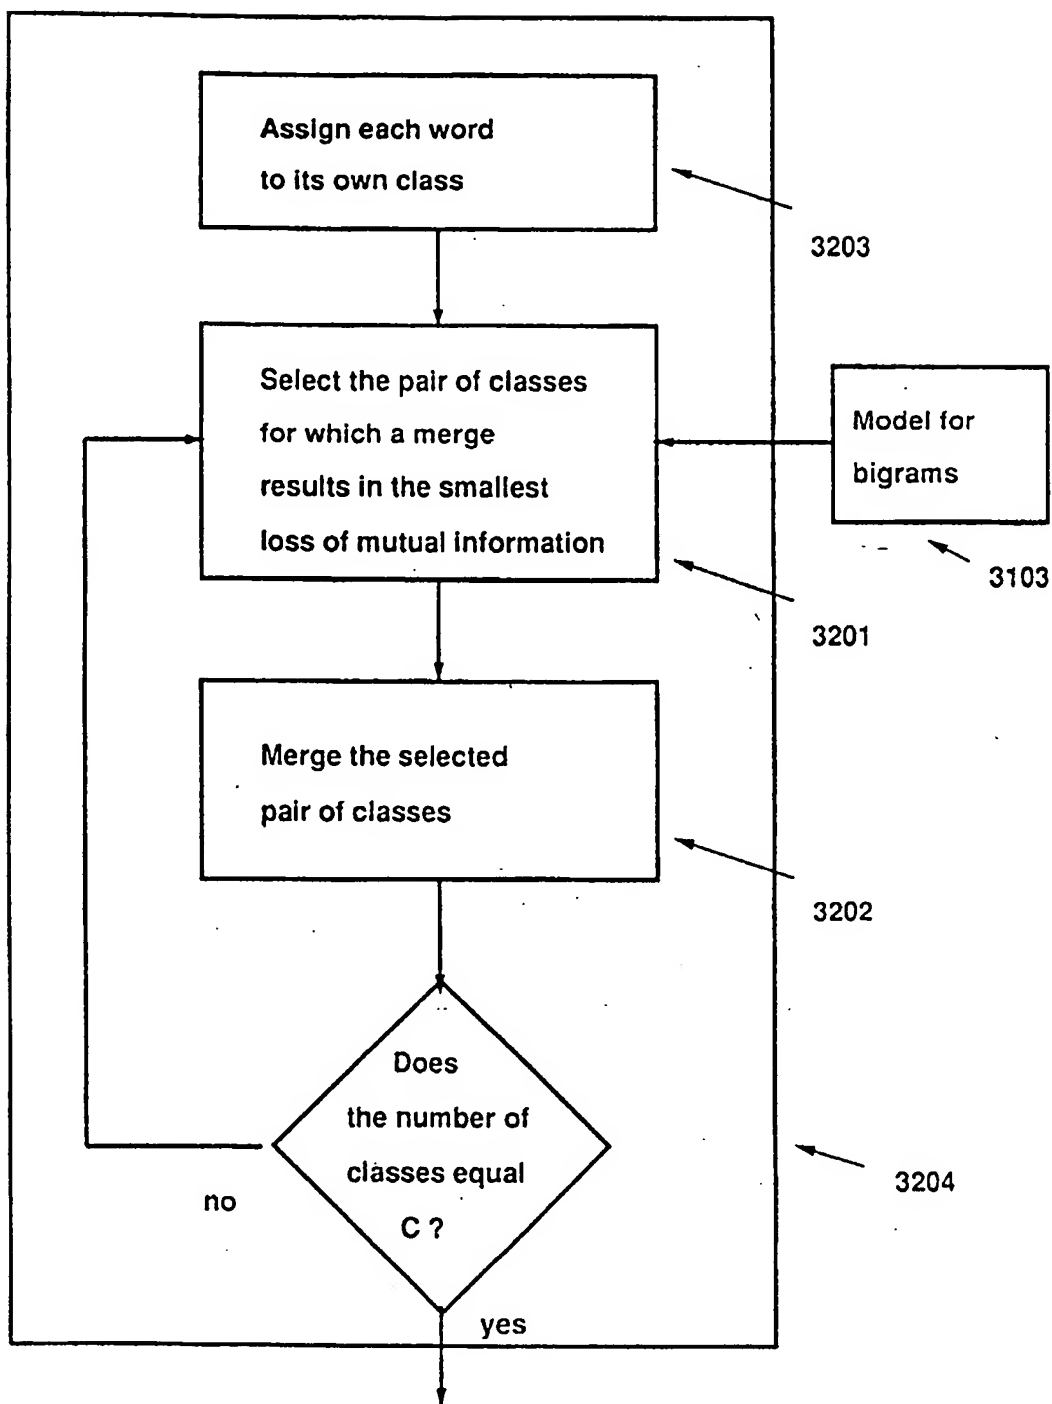


Figure 32




Figure 33




Figure 34




Figure 35




Figure 36




Figure 37






Figure 38


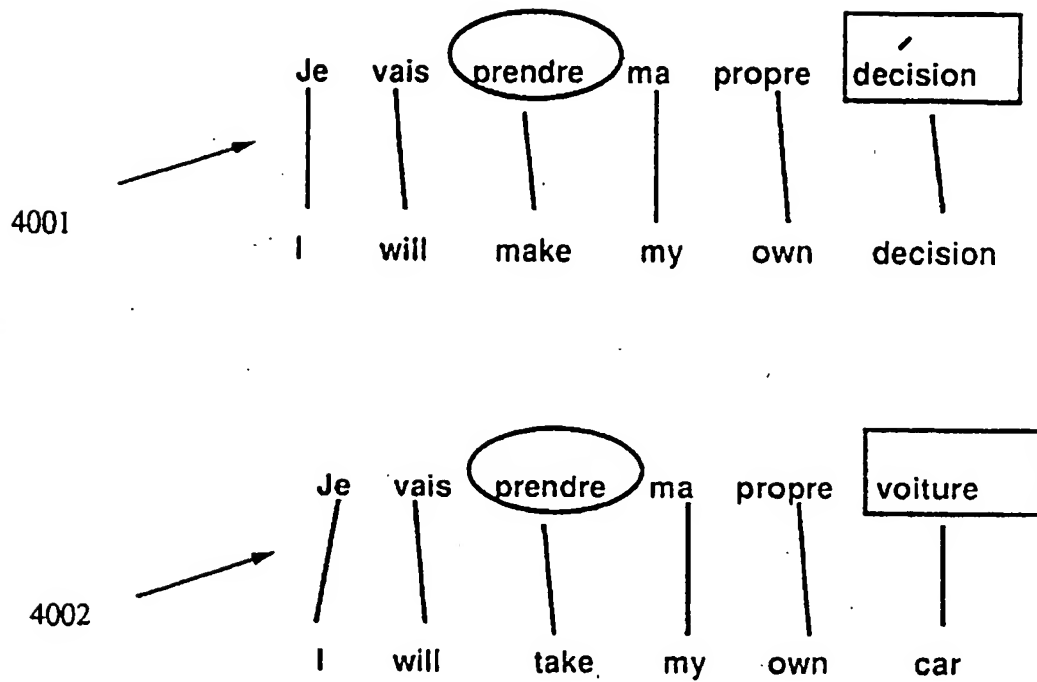


Figure 39



vocabulary word: prendre

informant site: first noun to the right

informant word:

first sequence: decision

second sequence: voiture

Figure 40

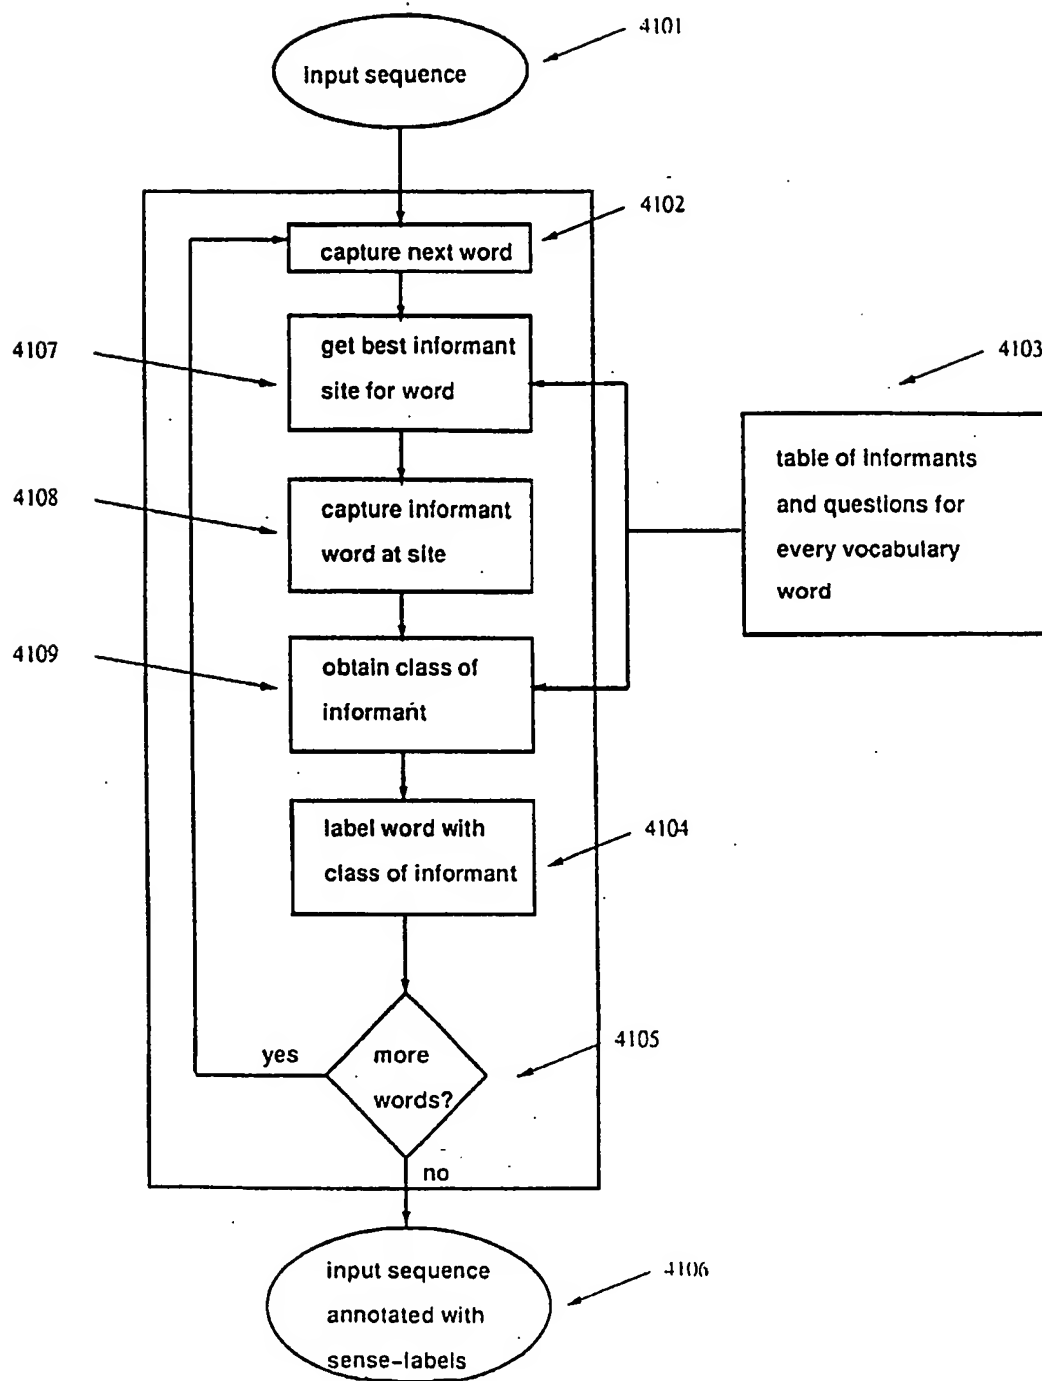


Figure 41

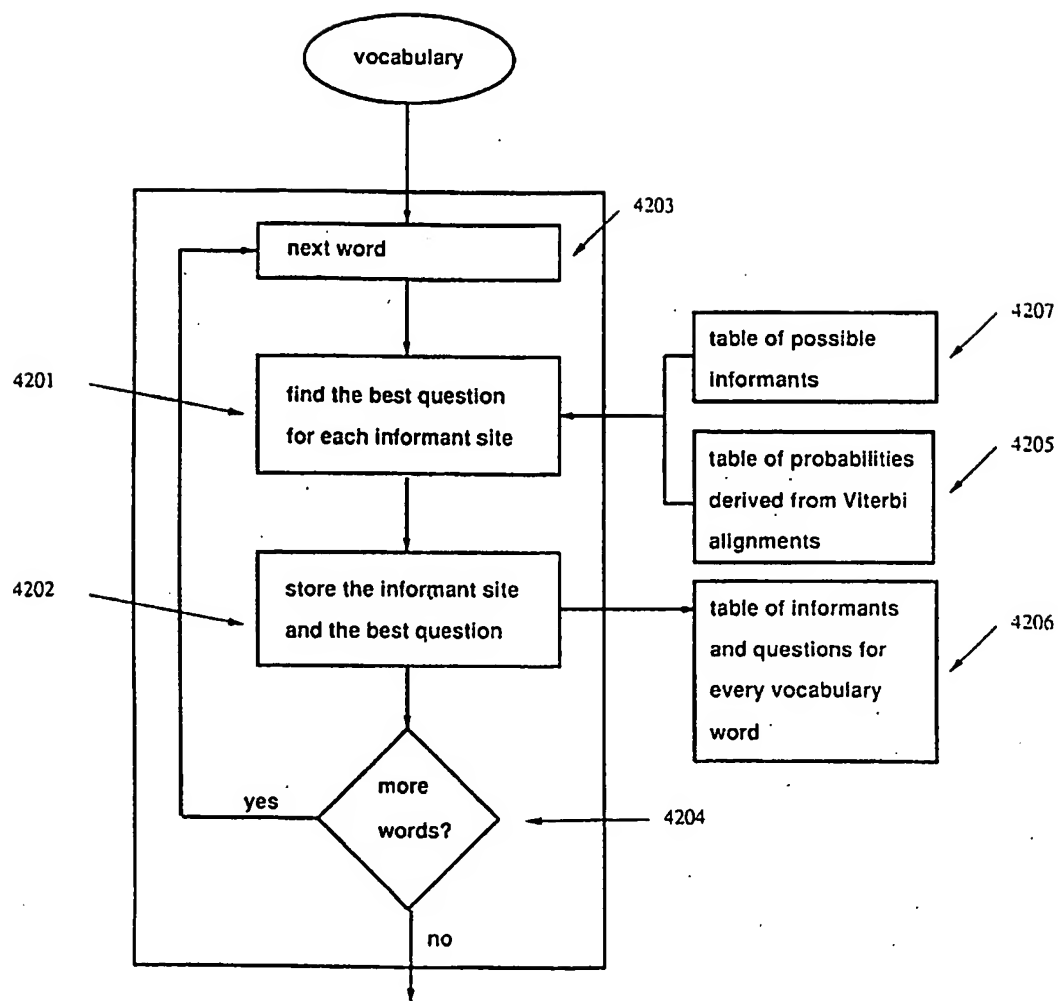


Figure 42

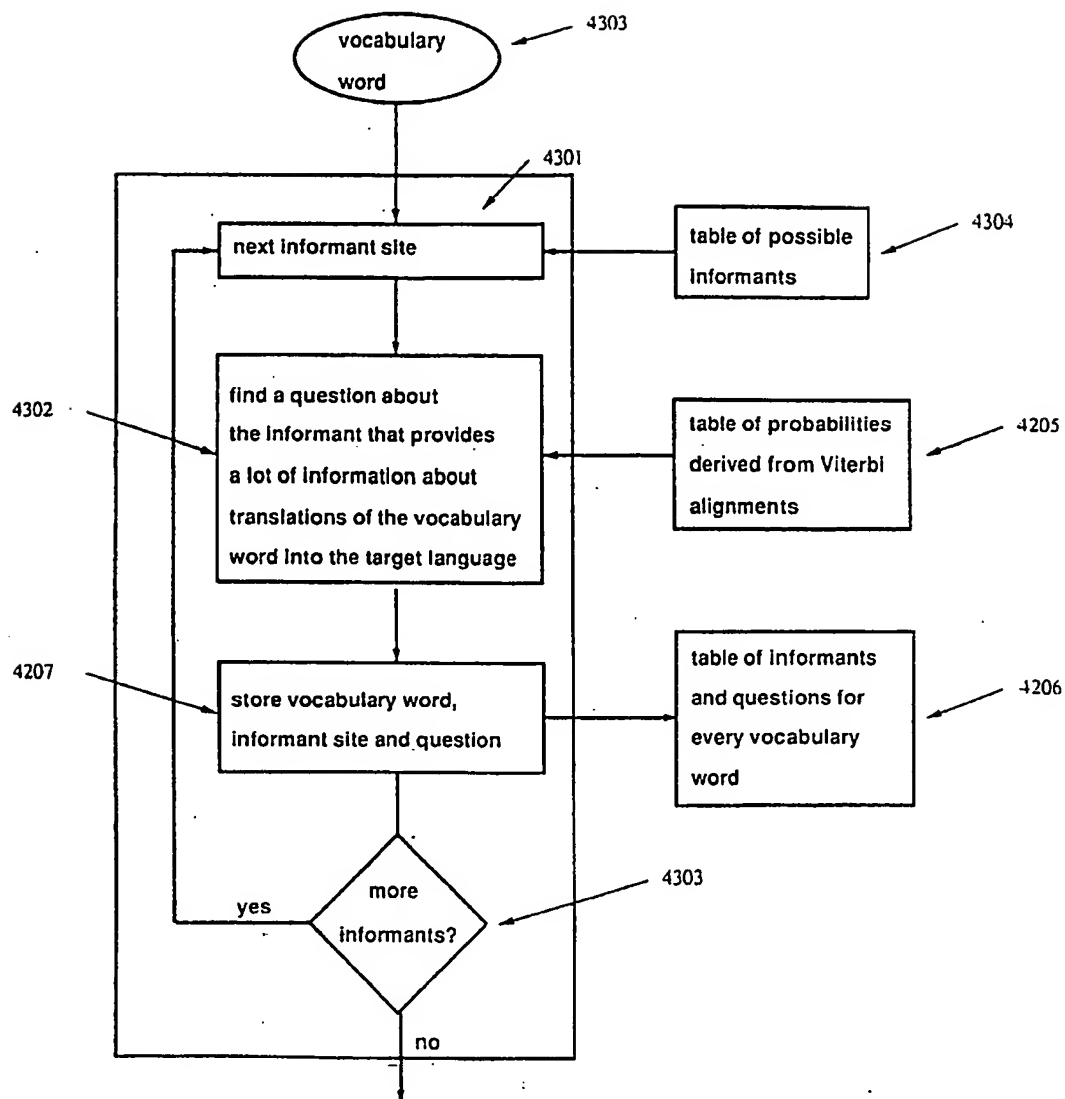


Figure 43




Figure 44




Figure 45



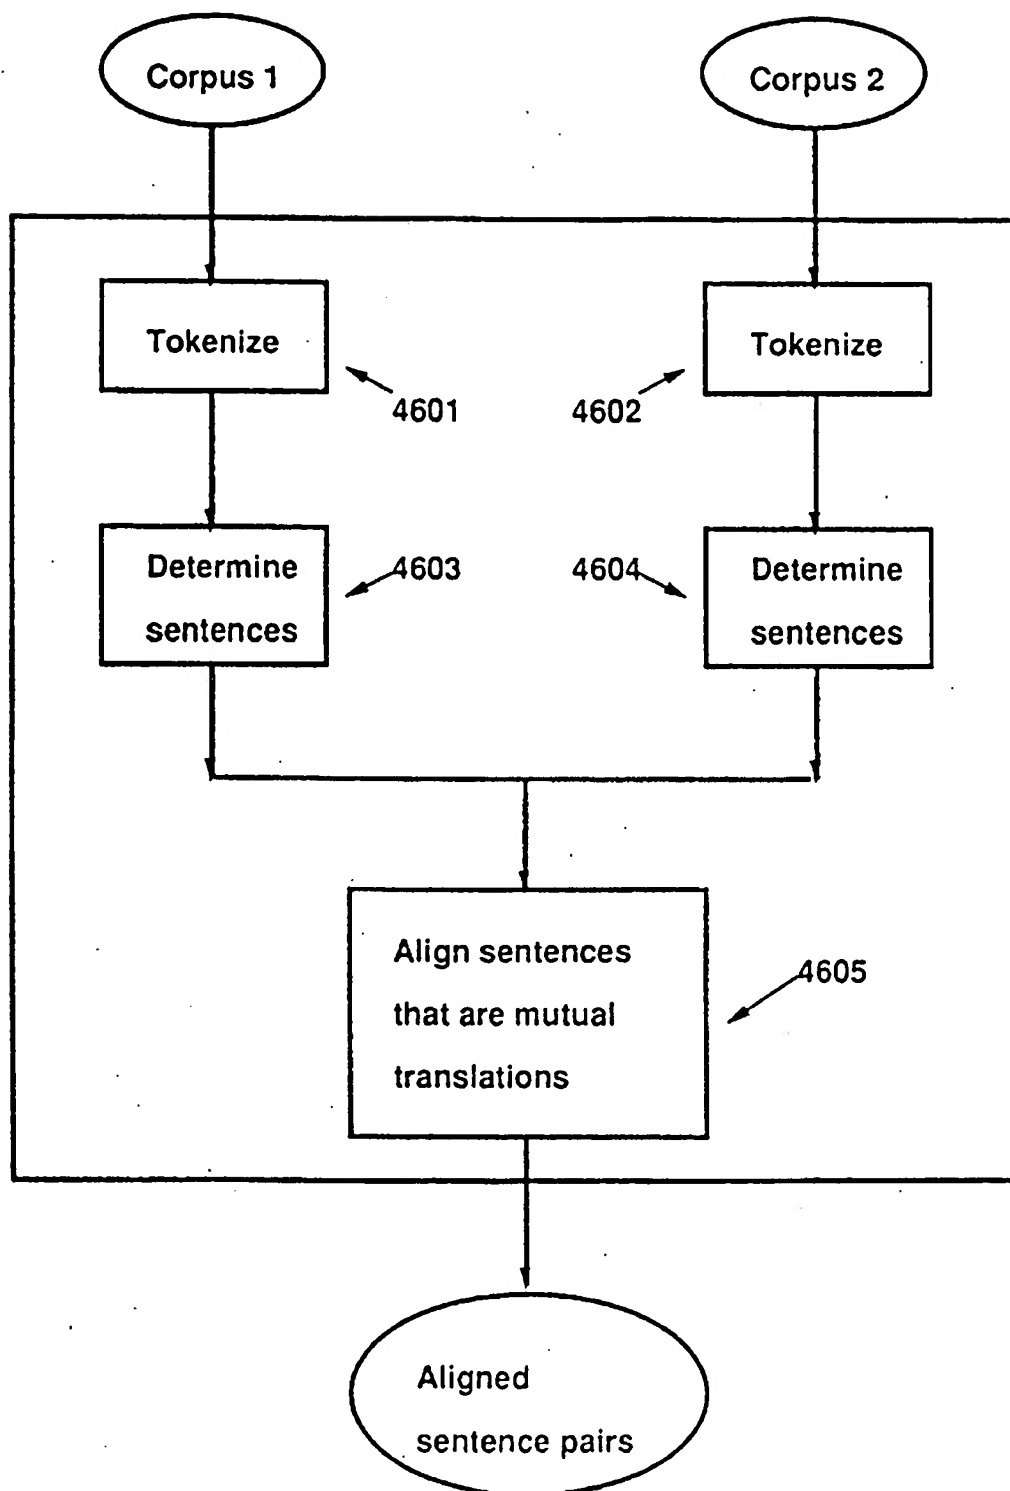


Figure 46




Figure 47

```

/*START_COMMENT* Beginning file = 048 1. \SCM{} Document = 048 101 H002-108
101 H002-108 script A *END_COMMENT*/      script A \ECM{}
.TB 029 060 090 099
.PL 060
.LL 120
.NF
The House met at 2 p.m.
.SP
*boMr. Donald MacInnis (Cape Breton
-East Richmond):*ro Mr. Speaker,
I rise on a question of privilege af-
fecting the rights and prerogatives
of parliamentary committees and one
which reflects on the word of two
ministers.
.SP
*boMr. Speaker: *roThe hon. member's
motion is proposed to the
House under the terms of Standing
Order 43. Is there unanimous consent?
.SP
*boSome hon. Members: *roAgreed.
s*itText*ro)
Question No. 17---*boMr. Mazankowski:
*ro
1. For the period April 1, 1973 to
January 31, 1974, what amount of
money was expended on the operation
and maintenance of the Prime
Minister's residence at Harrington
Lake, Quebec?
.SP
(1415)
s*itLater:*ro)
.SP
*boMr. Cossitt:*ro Mr. Speaker, I rise
on a point of order to ask for
clarification by the parliamentary
secretary.
2. The House met at 2 p.m.
3. \SCM{} Paragraph \ECM{}
4. \SCM{} Author = Mr. Donald MacInnis
(Cape Breton-East Richmond) \ECM{}
5. Mr. Speaker, I rise on a question of
privilege affecting the rights and
prerogatives of parliamentary
committees and one which reflects on
the word of two ministers.
21. \SCM{} Paragraph \ECM{}
22. \SCM{} Author = Mr. Speaker \ECM{}
23. The hon. member's motion is proposed
to the House under the terms of
Standing Order 43.
44. Is there unanimous consent?
45. \SCM{} Paragraph \ECM{}
46. \SCM{} Author = Some hon. Members
\ECM{}
47. Agreed.
61. \SCM{} Source = Text \ECM{}
62. \SCM{} Question = 17 \ECM{}
63. \SCM{} Author = Mr. Mazankowski
\ECM{}
64. 1.
65. For the period April 1, 1973 to
January 31, 1974, what amount of
money was expended on the operation
and maintenance of the Prime
Minister's residence at Harrington
Lake, Quebec?
66. \SCM{} Paragraph \ECM{}
81. \SCM{} Time = (1415) \ECM{}
82. \SCM{} Time = Later \ECM{}
83. \SCM{} Paragraph \ECM{}
84. \SCM{} Author = Mr. Cossitt \ECM{}
85. Mr. Speaker, I rise on a point of
order to ask for clarification by
the parliamentary secretary.

```

Figure 48




Figure 49



Figure 50

<i>Bead</i>	<i>Text</i>
<i>e</i>	one English sentence
<i>f</i>	one French sentence
<i>ef</i>	one English and one French sentence
<i>eeff</i>	two English and one French sentence
<i>eff</i>	one English and two French sentences
<i>¶</i>	one English paragraph
<i>¶<sub>f</sub></i>	one French paragraph
<i>¶, ¶<sub>f</sub></i>	one English and one French paragraph

Figure 51




Figure 52




Figure 53

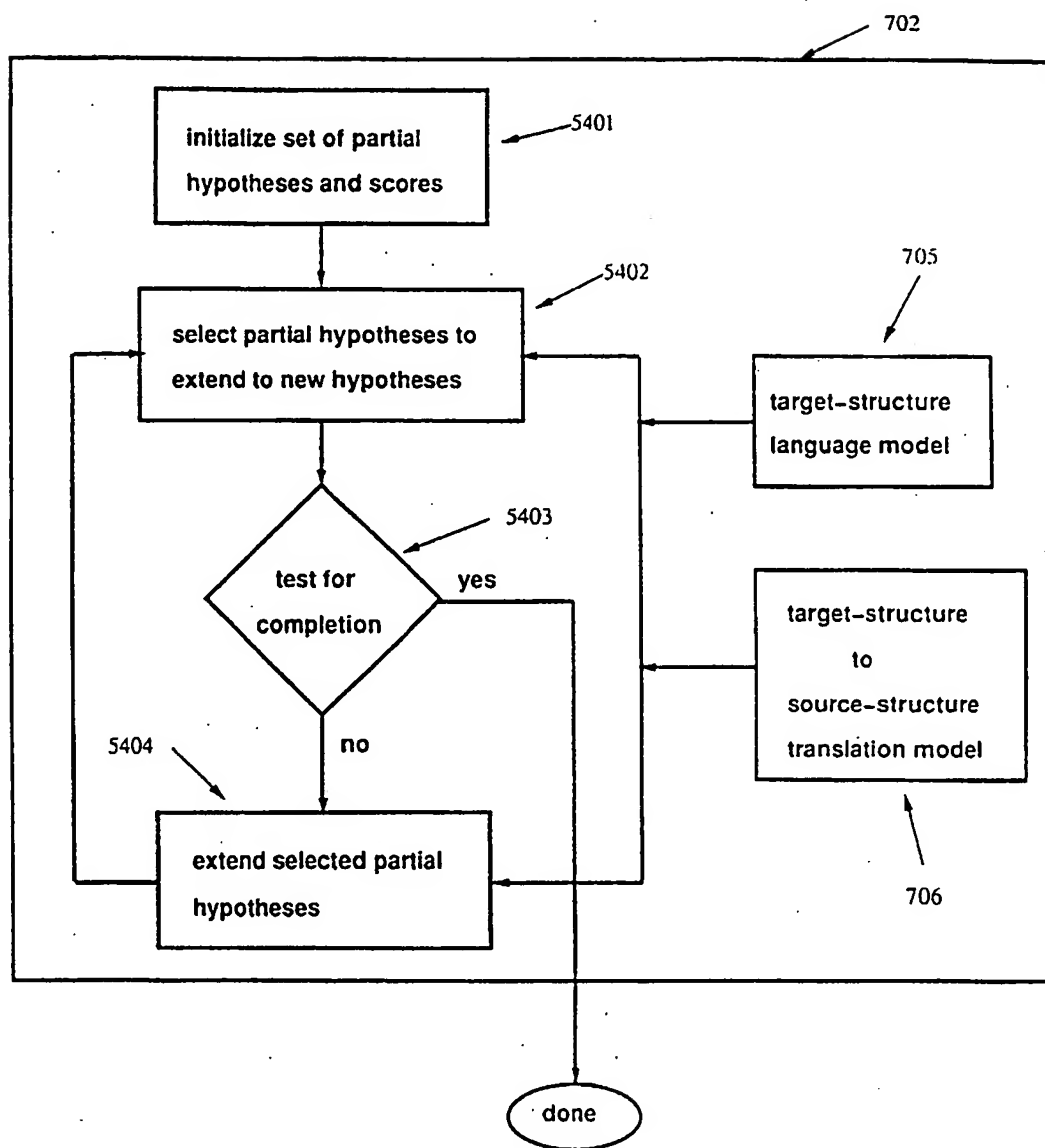


Figure 54

La jeune fille a réveillé sa mère

transduce  
source text

la jeune fille V\_past\_3s réveiller sa mère

Figure 55

the

la jeune fille V\_past\_3s réveiller sa mère

Figure 56

mother

la jeune fille V\_past\_3s réveiller sa mère

Figure 57






Figure 58




Figure 59




Figure 60




Figure 61




Figure 62




Figure 63




Figure 64

yes ,  
oui , oui

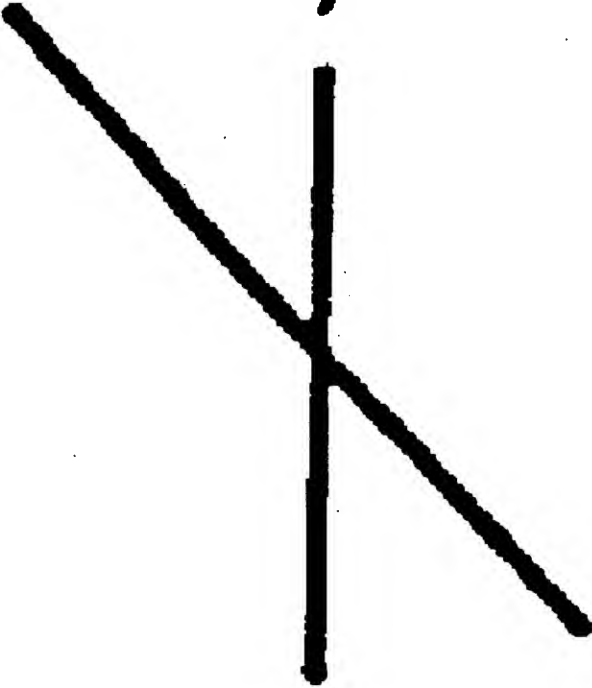
A large, thick, black handwritten 'X' mark is drawn across the text. It starts from the top left, crosses the middle, and ends at the bottom right, effectively crossing out the words 'yes' and 'oui'.

Figure 65

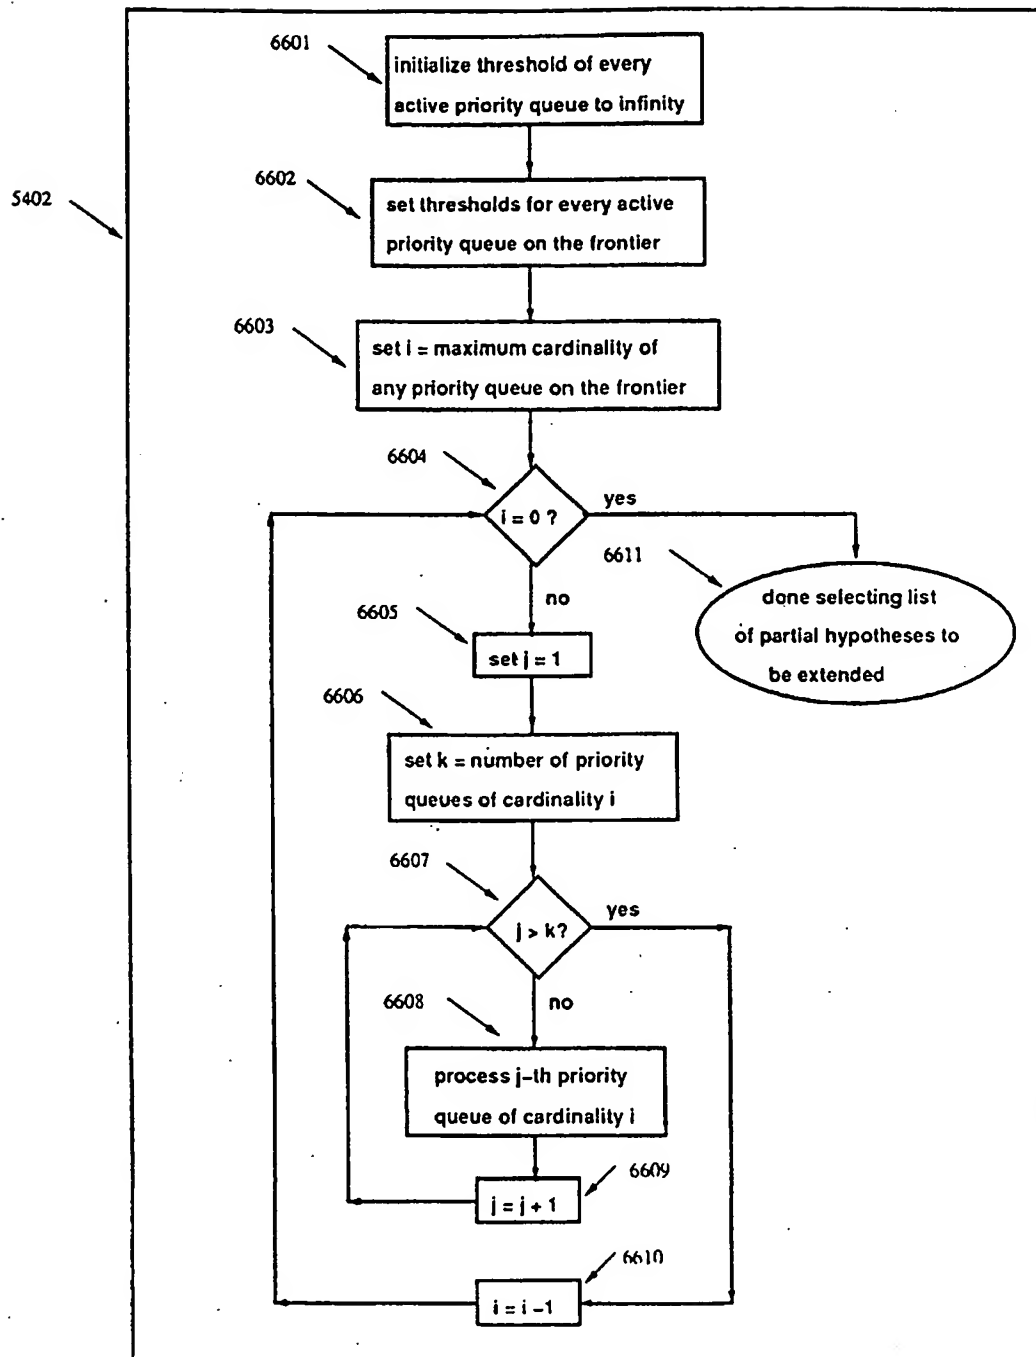


Figure 66

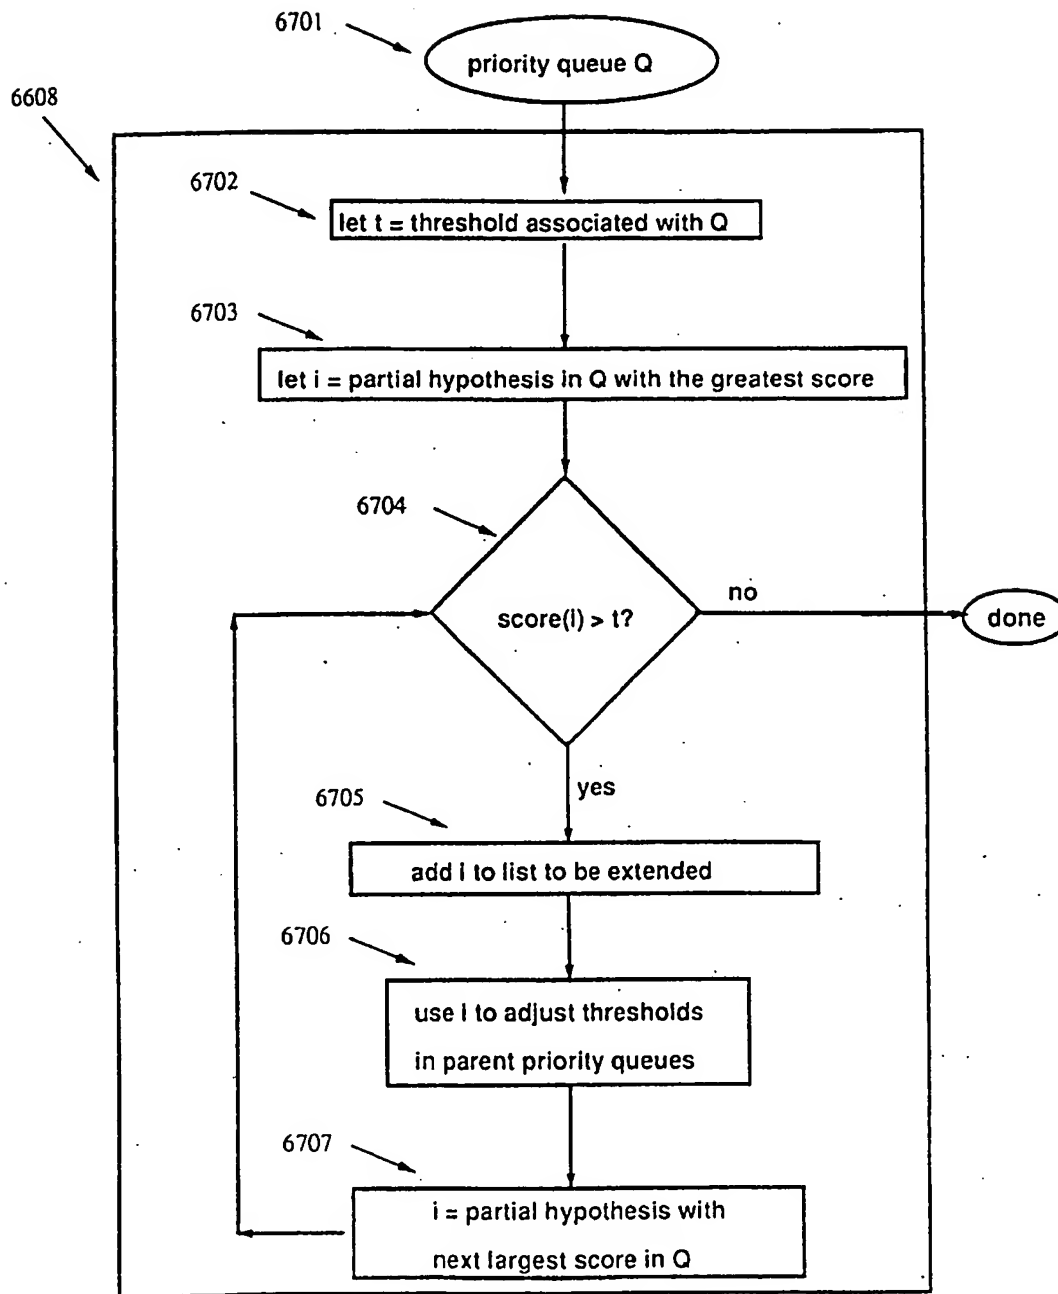


Figure 67

```
1 procedure: extend_partial_hypotheses_on_list;
2   do for i = 1 to the number of partial hypotheses on list;
3     let h be the i-th partial hypothesis to be extended;
4     do for p = 1 to the number of positions in the source structure.
5       if p is not already aligned in h then
6         extend_h_by_accounting_for_source_morpheme_in_position p;
7       end of do for p;
8     end of do for i;
9   end of procedure extend_partial_hypotheses_on_list;

10 procedure: extend_h_by_accounting_for_source_morpheme_in_position_p;
11   if h is an open partial hypothesis then do;
12     let q be the position of the open target morpheme in h;
13     extend_open_h_by_connecting_p_to_q_and_keep_h_open;
14     extend_open_h_by_connecting_p_to_q_and_close_h;
15   end of if h;
16   if h is not an open partial hypothesis then do;
17     let s be the source morpheme in position p;
18     do j = 1 to the number of target morpheme translations of s;
19       let t be the j-th target morpheme translation of s;
20       create_open_extension_of_h_by_connecting_p_to_t;
21       create_closed_extension_of_h_by_connecting_p_to_t;
22       create_extension_of_h_by_connecting_p_to_null_target_morpheme;
23       create_list_of_target_morphemes_to_be_inserted_before_t;
24       do k = 1 to number of target morphemes on listed to be inserted;
25         let t1 be the k-th target morpheme to be inserted;
26         create_open_extenstion_of_by_connecting_p_to_t1_t;
27         create_closed_extenstion_of_by_connecting_p_to_t1_t;
28       end of do k;
29     end of do j;
30   end of if h;
31 end of procedure extend_h_by_accounting_for_source_morpheme_in_position_p;
```

Figure 68

## METHOD AND SYSTEM FOR NATURAL LANGUAGE TRANSLATION

### TECHNICAL FIELD

The present invention relates to machine translation in general, and, in particular to methods of performing machine translation which use statistical models of the translation process.

### BACKGROUND ART

There has long been a desire to have machines capable of translating text from one language into text in another language. Such machines would make it much easier for humans who speak different languages to communicate with one another. The present invention uses statistical techniques to attack the problem of machine translation. Related statistical techniques have long been used in the field of automatic speech recognition. The background for the invention is thus in two different fields: automatic speech recognition and machine translation.

The central problem in speech recognition is to recover from an acoustic signal the word sequence which gave rise to it. Prior to 1970, most speech recognition systems were built around a set of hand-written rules of syntax, semantics and acoustic-phonetics. To construct such a system it is necessary to firstly discover a set of linguistic rules that can account for the vast complexity of language, and, secondly to construct a coherent framework in which these rules can be assembled to recognize speech. Both of these problems proved insurmountable. It proved too difficult to write down by hand a set of rules that adequately covered the vast scope of natural language and to construct by hand a set of weights, priorities and if-then statements that can regulate interactions among its many facets.

This impasse was overcome in the early 1970's with the introduction of statistical techniques to speech recognition. In the statistical approach, linguistic rules are extracted automatically using statistical techniques from large databases of speech and text. Different types of linguistic information are combined via the formal laws of probability theory. Today, almost all speech recognition systems are based on statistical techniques.

Speech recognition has benefited by using statistical language models which exploit the fact that not all word sequences occur naturally with equal probability. One simple model is the trigram model of English, in which it is assumed that the probability that a word will be spoken depends only on the previous two words that have been spoken. Although trigram models are simple-minded, they have proven extremely powerful in their ability to predict words as they occur in natural language, and in their ability to improve the performance of natural-language speech recognition. In recent years more sophisticated language models based on probabilistic decision-trees, stochastic context-free grammars and automatically discovered classes of words have also been used.

In the early days of speech recognition, acoustic models were created by linguistic experts, who expressed their knowledge of acoustic-phonetic rules in programs which analyzed an input speech signal and produced as output a sequence of phonemes. It was thought to be a simple matter to decode a word sequence from a sequence of phonemes. It turns out, however, to be a very difficult job to determine an accurate phoneme sequence from a speech signal. Although

human experts certainly do exist, it has proven extremely difficult to formalize their knowledge. In the alternative statistical approach, statistical models, most predominantly hidden Markov models, capable of learning acoustic-phonetic knowledge from samples of speech are employed.

The present approaches to machine translation are similar in their reliance on hand-written rules to the approaches to speech recognition twenty years ago. Roughly speaking, the present approaches to machine translation can be classified into one of three categories: direct, interlingual, and transfer. In the direct approach a series of deterministic linguistic transformations is performed. These transformations directly convert a passage of source text into a passage of target text. In the transfer approach, translation is performed in three stages: analysis, transfer, and synthesis. The analysis stage produces a structural representation which captures various relationships between syntactic and semantic aspects of the source text. In the transfer stage, this structural representation of the source text is then transferred by a series of deterministic rules into a structural representation of the target text. Finally, in the synthesis stage, target text is synthesized from the structural representation of the target text. The interlingual approach to translation is similar to the transfer approach except that in the interlingual approach an internal structural representation that is language independent is used. Translation in the interlingual approach takes place in two stages, the first analyzes the source text into this language-independent interlingual representation, the second synthesizes the target text from the interlingual representation. All these approaches use hand-written deterministic rules.

Statistical techniques in speech recognition provide two advantages over the rule-based approach. First, they provide means for automatically extracting information from large bodies of acoustic and textual data, and second, they provide, via the formal rules of probability theory, a systematic way of combining information acquired from different sources. The problem of machine translation between natural languages is an entirely different problem than that of speech recognition. In particular, the main area of research in speech recognition, acoustic modeling, has no place in machine translation. Machine translation does face the difficult problem of coping with the complexities of natural language. It is natural to wonder whether this problem won't also yield to an attack by statistical methods, much as the problem of coping with the complexities of natural speech has been yielding to such an attack. Although the statistical models needed would be of a very different nature, the principles of acquiring rules automatically and combining them in a mathematically principled fashion might apply as well to machine translation as they have to speech recognition.

### DISCLOSURE OF THE INVENTION

The present invention is directed to a system and method for translating source text from a first language to target text in a second language different from the first language. The source text is first received and stored in a first memory buffer. Zero or more user defined criteria pertaining to the source text are also received and stored by the system. The user defined criteria are used by various subsystems to bound the target text.

The source text is then transduced into one or more intermediate source structures which may be constrained by any of the user defined criteria. One or more target hypotheses are then generated. Each of the target hypotheses comprise an intermediate target structure of text selected from the second language. The intermediate target structures may

also be constrained by any of the user defined criteria.

A target-structure language model is used to estimate a first score which is proportional to the probability of occurrence of each intermediate target-structure of text associated with the target hypotheses. A target structure-to-source-structure translation model is used to estimate a second score which is proportional to the probability that the intermediate target-structure of text associated with the target hypotheses will translate into the intermediate source-structure of text. For each target hypothesis, the first and second scores are combined to produce a target hypothesis match score.

The intermediate target-structures of text are then transduced into one or more transformed target hypothesis of text in the second language. This translation step may also be constrained by any of the user defined criteria.

Finally, one or more of the transformed target hypotheses is displayed to the user according to its associated match score and the user defined criteria. Alternatively, one or more of the transformed target hypotheses may be stored in memory or otherwise made available for future access.

The intermediate source and target structures are transduced by arranging and rearranging, respectively, elements of the source and target text according to one or more of a lexical substitution, part-of-speech assignment, a morphological analysis, and a syntactic analysis.

The intermediate target structures may be expressed as an ordered sequence of morphological units, and the first score probability may be obtained by multiplying the conditional probabilities of each morphological unit within an intermediate target structure given the occurrence of previous morphological units within the intermediate target structure. In another embodiment, the conditional probability of each unit of each of the intermediate target structure may depend only on a fixed number of preceding units within the intermediate target structure.

The intermediate source structure transducing step may further comprise the step of tokenizing the source text by identifying individual words and word separators and arranging the words and the word separators in a sequence. Moreover, the intermediate source structure transducing step may further comprise the step of determining the case of the words using a case transducer, wherein the case transducer assigns each word a token and a case pattern. The case pattern specifies the case of each letter of the word. Each case pattern is evaluated to determine the true case pattern of the word.

#### BRIEF DESCRIPTION OF DRAWINGS

The above stated and other aspects of the present invention will become more evident upon reading the following description of preferred embodiments in conjunction with the accompanying drawings, in which:

FIG. 1 is a schematic block diagram of a simplified French to English translation system;

FIG. 2 is a schematic block diagram of a simplified French to English translation system which incorporates analysis and synthesis;

FIG. 3 is a schematic block diagram illustrating a manner in which statistical transfer can be incorporated into a translation system based on an interlingua;

FIG. 4 is a schematic flow diagram of a basic system operating in batch mode;

FIG. 5 is a schematic flow diagram of a basic system operating in user-aided mode;

FIG. 6 is a schematic representation of a work station environment through which a user interfaces with a user-aided system;

FIG. 7 is a schematic flow diagram of a translation component of a batch system;

FIG. 8 is a schematic flow diagram illustrating a user's interaction with a human-aided system;

FIG. 9 is a schematic flow diagram of a translation component of a human-aided system;

FIG. 10 is a schematic block diagram of a basic hardware components used in a preferred embodiment of the present invention;

FIG. 11 is a schematic flow diagram of a source transducer;

FIG. 12 is a schematic flow diagram of a token-to-word transducer;

FIG. 13 is a schematic flow diagram of a syntactic transducer for English;

FIG. 14 is a schematic flow diagram of a syntactic transducer for French.

FIG. 15 is an example of a syntactic transduction;

FIG. 16 is a schematic flow diagram illustrating the operation of a finite-state transducer;

FIG. 17 is a Table of Patterns for a finite-state pattern matcher employed in some source and target transducers;

FIG. 18 is a Table of Actions for an action processing module employed in some source and target transducers.

FIG. 19 is a schematic flow diagram of a target transducer;

FIG. 20 is a schematic block diagram of a target structure language model

FIG. 21 is a schematic block diagram of a target structure to source structure translation model;

FIG. 22 is a simple example of an alignment between a source structure and a target structure;

FIG. 23 is another example of an alignment;

FIG. 24 is a schematic flow diagram of showing an embodiment of a target structure to source structure translation model;

FIG. 25 is a schematic block diagram of a detailed translation model;

FIG. 26 is a schematic flow diagram of an embodiment of a detailed translation model;

FIG. 27 is a schematic flow diagram of a method for estimating parameters for progression of gradually more sophisticated translation models;

FIG. 28 is a schematic flow diagram of a method for iteratively improving parameter estimates for a model;

FIG. 29 is a schematic flow diagram of a method for using parameters values for one model to obtain good initial estimates of parameters for another model;

FIG. 30 shows some sample subtrees from a tree constructed using a clustering method;

FIG. 31 is a schematic block diagram of a method for partitioning a vocabulary into classes;

FIG. 32 is a schematic flow diagram of a method for partitioning a vocabulary into classes;

FIG. 33 is an alignment with independent English words;

FIG. 34 is an alignment with independent French words;

FIG. 35 is a general alignment;

FIG. 36 shows two tableaux for one alignment;

FIG. 37 shows the best alignment out of  $1.9 \times 10^{23}$  alignments;



FIG. 38 shows the best of  $8.4 \times 10^{29}$  possible alignments;  
FIG. 39 shows the best of  $5.6 \times 10^{31}$  alignments;

FIG. 40 is an example of informants and informant sites;

FIG. 41 is a schematic flow diagram illustrating the operation of a sense-labeling transducer;

FIG. 42 is a schematic flow diagram of a module that determines good questions about informants for each vocabulary word;

FIG. 43 is a schematic flow diagram of a module that determines a good question about each informant of a vocabulary word;

FIG. 44 is a schematic flow diagram of a method for determining a good question about an informant;

FIG. 45 is a schematic diagram of a process by which word-by-word correspondences are extracted from a bilingual corpus;

FIG. 46 is a schematic flow diagram of a method for aligning sentences in parallel corpora;

FIG. 47 is a schematic flow diagram of the basic step of a method for aligning sentences;

FIG. 48 depicts a sample of text before and after textual cleanup and sentence detection;

FIG. 49 shows an example of a division of aligned corpora into beads;

FIG. 50 shows a finite state model for generating beads;

FIG. 51 shows the beads that are allowed by the model;

FIG. 52 is a histogram of French sentence lengths;

FIG. 53 is a histogram of English sentence lengths;

FIG. 54 is a schematic flow diagram of the hypothesis search component of the system;

FIG. 55 is an example of a source sentence being transduced to a sequence of morphemes;

FIG. 56 is an example of a partial hypothesis which results from an extension by the target morpheme the;

FIG. 57 is an example of a partial hypothesis which results from an extension by the target morpheme mother;

FIG. 58 is an example of a partial hypothesis which results from an extension with an open target morpheme;

FIG. 59 is an example of a partial hypothesis which results from an extension in which an open morpheme is closed;

FIG. 60 is an example of a partial hypothesis which results from an extension in which an open morpheme is kept open;

FIG. 61 is an example of a partial hypothesis which results from an extension by the target morpheme to\_wake;

FIG. 62 is an example of a partial hypothesis which results from an extension by the pair of target morphemes up and to\_wake;

FIG. 63 is an example of a partial hypothesis which results from an extension by the pair of target morphemes up and to\_wake in which to\_wake is open;

FIG. 64 is an example of a subset lattice;

FIG. 65 is an example of a partial hypothesis that is stored in the priority queue {2,3};

FIG. 66 is a schematic flow diagram of the process by which partial hypotheses are selected for extension;

FIG. 67 is a schematic flow diagram of the method by which the partial hypotheses on a priority queue are processed in the selection for extension step;

FIG. 68 contains pseudocode describing the method for extending hypotheses.

## BEST MODE FOR CARRYING OUT THE INVENTION

### Contents

1	Introduction
2	Translation Systems
3	Source Transducers
3.1	Overview
3.2	Components
3.2.1	Tokenizing Transducers
3.2.2	Token-Word Transducers
3.2.3	True-Case Transducers
3.2.4	Specialized Transformation Transducers
3.2.5	Part-of-Speech Labeling Transducers
3.2.6	Syntactic Transformation Transducers
3.2.7	Syntactic Transducers for English
3.2.8	Syntactic Transducers for French
3.2.9	Morphological Transducers
3.2.10	Sense-Labeling Transducers
3.3	Source-Transducers with Constraints
4	Finite-State Transducers
5	Target Transducers
5.1	Overview
5.2	Components
5.3	Target Transducers with Constraints
6	Target Language Model
6.1	Perplexity
6.2	n-gram Language models
6.3	Simple n-gram models
6.4	Smoothing
6.5	n-gram Class Models
7	Classes
7.1	Maximum Mutual Information Clustering
7.2	A Clustering Method
7.3	An Alternate Clustering Method
7.4	Improving Classes
7.5	Examples
7.6	A Method for Constructing Similarity Trees
8	Overview of Translation Models and Parameter Estimation
8.1	Translation Models
8.1.1	Training
8.1.2	Hidden Alignment Models
8.2	An example of a detailed translation model
8.3	Iterative Improvement
8.3.1	Relative Objective Function
8.3.2	Improving Parameter Values
8.3.3	Going From One Model to Another
8.3.4	Parameter Reestimation Formulae
8.3.5	Approximate and Viterbi Parameter Estimation
8.4	Five Models
9	Detailed Description of Translation Models and Parameter Estimation
9.1	Notation
9.2	Translations
9.3	Alignments
9.4	Model 1
9.5	Model 2
9.6	Intermodel interlude
9.7	Model 3
9.8	Deficiency
9.9	Model 4
9.10	Model 5
9.11	Results
9.12	Better Translation Models
9.12.1	Deficiency
9.12.2	Multi-word Notions
10	Mathematical Summary of Translation Models

- 10.1 Summary of Notation
- 10.2 Model 1
- 10.3 Model 2
- 10.4 Model 3
- 10.5 Model 4
- 10.6 Model 5
- 11 Sense Disambiguation
  - 11.1 Introduction
  - 11.2 Design of a Sense-Labeling Transducer
  - 11.3 Constructing a table of informants and questions
  - 11.4 Mathematics of Constructing Questions
    - 11.4.1 Statistical Translation with Transductions
    - 11.4.2 Sense-Labeling in Statistical Translation
    - 11.4.3 The Viterbi Approximation
    - 11.4.4 Cross Entropy
  - 11.5 Selecting Questions
    - 11.5.1 Source Questions
    - 11.5.2 Target Questions
  - 11.6 Generalizations
- 12 Aligning Sentences
  - 12.1 Overview
  - 12.2 Tokenization and Sentence Detection
  - 12.3 Selecting Anchor Points
  - 12.4 Aligning Major Anchors
  - 12.5 Discarding Sections
  - 12.6 Aligning Sentences
  - 12.7 Ignoring Anchors
  - 12.8 Results for the Hansard Example
- 13 Aligning Bilingual Corpora
- 14 Hypothesis Search—Steps 702 and 902
  - 14.1 Overview of Hypothesis Search
  - 14.2 Hypothesis Extension 5404
    - 14.2.1 Types of Hypothesis Extension
  - 14.3 Selection of Hypotheses to Extend 5402

# 1 Introduction

The inventions described in this specification include apparatuses for translating text in a source language to text in a target language. The inventions employ a number of different components. There are various embodiments of each component, and various embodiments of the apparatuses in which the components are connected together in different ways.

To introduce some of the components, a number of simplified embodiments of a translation apparatus will be described in this introductory section. The components themselves will be described in detail in subsequent sections. More sophisticated embodiments of the apparatuses, including the preferred embodiments, will also be described in subsequent sections.

The simplified embodiments translate from French to English. These languages are chosen only for the purpose of example. The simplified embodiments, as well as the preferred embodiments, can easily be adapted to other natural languages, including but not limited to, Chinese, Danish, Dutch, English, French, German, Greek, Italian, Japanese, Portuguese, and Spanish, as well as to artificial languages including but not limited to database query languages such as SQL, and the programming languages such as LISP. For example, in one embodiment queries in English might be translated into an artificial database query language.

A simplified embodiment 104 is depicted in FIG. 1. It includes three components:

1. A language model 101 which assigns a probability or score  $P(E)$  to any portion of English text  $E$ ;
2. A translation model 102 which assigns a conditional probability or score  $P(F|E)$  to any portion of French text  $F$  given any portion of English text  $E$ ; and

3. A decoder 103 which given a portion of French text  $F$  finds a number of portions of English text  $E$ , each of which has a large combined probability or score

$$P(E,F)=P(E)P(F|E). \quad (1)$$

Language models are described in Section 6, and translation models are described in Sections 8, 9, and 10. An embodiment of the decoder 103 is described in detail in Section 14.

A shortcoming of the simplified architecture depicted in FIG. 1 is that the language model 101 and the translation model 102 deal directly with unanalyzed text. The linguistic information in a portion of text and the relationships between translated portions of text is complex, involving linguistic phenomena of a global nature. However, the models 101 and 102 must be relatively simple so that their parameters can be reliably estimated from a manageable amount of training data. In particular, they are restricted to the description of local structure.

This difficulty is addressed by integrating the architecture of FIG. 1 with the traditional machine translation architecture of analysis, transfer, and synthesis. The result is the simplified embodiment depicted in FIG. 2. It includes six components:

1. a French transduction component 201 in which a portion of French text  $F$  is encoded into an intermediate structure  $F'$ ;
2. a statistical transfer component 202 in which  $F'$  is translated to a corresponding intermediate English structure  $E'$ ;
3. an English transduction component 203 in which a portion of English text  $E$  is reconstructed from  $E'$ ;

This statistical transfer component 202 is similar in function to the whole of the simplified embodiment 104, except that it translates between intermediate structures instead of translating directly between portions of text. The building blocks of these intermediate structures may be words, morphemes, syntactic markers, semantic markers, or other units of linguistic structure. In descriptions of aspects of the statistical transfer component, for reasons of exposition, these building blocks will sometimes be referred to as words or morphemes, morphs, or morphological units. It should be understood, however, that the unless stated otherwise the same descriptions of the transfer component apply also to other units of linguistic structure.

The statistical transfer component incorporates

4. an English structure language model 204 which assigns a probability or score  $P(E')$  to any intermediate structure  $E'$ ;
5. an English structure to French structure translation model 205 which assigns a conditional probability or score  $P(F'|E')$  to any intermediate structure  $F'$  given any intermediate structure  $E'$ ; and
6. a decoder 206 which given a French structure  $F'$  finds a number of English structures  $E'$ , each of which has a large combined probability or score

$$P(E',F')=P(E')P(F'|E'). \quad (2)$$

The role of the French transduction component 201 and the English transduction component 203 is to facilitate the task of the statistical transfer component 202. The intermediate structures  $E'$  and  $F'$  encode global linguistic facts about  $E$  and  $F$  in a local form. As a result, the language and translation models 204 and 205 for  $E'$  and  $F'$  will be more accurate than the corresponding models for  $E$  and  $F$ .

The French transduction component 201 and the English transduction component 203 can be composed of a sequence of successive transformation in which F' is incrementally constructed from F, and E is incrementally recovered from E'. For example, one transformation of the French transduction component 201 might label each word of F with the part of speech, e.g. noun, verb, or adjective, that it assumes in F. The part of speech assumed by a word depends on other words of F, but in F' it is encoded locally in the label. Another transformation might label each word of F with a cross-lingual sense label. The label of a word is designed to elucidate the English translation of the word.

Various embodiments of French transduction components are described in Section 3, and various embodiments of English transduction components are described in Section 5. Embodiments of a cross-lingual sense labeler and methods for sense-labeling are described in Section 11.

FIG. 2 illustrates the manner by which statistical transfer is incorporated into the traditional architecture of analysis, transfer, and synthesis. However, in other embodiments statistical transfer can be incorporated into other architectures. FIG. 3, for example, shows how two separate statistical transfer components can be incorporated into a French-to-English translation system which uses an interlingua. The French transducer 302 together with the statistical transfer module 304 translate French text 301 into an interlingual representations 305. The other statistical transfer module 306 together with the English transduction module 308 translate the interlingual representations 305 into English text 309.

The language models and translation models are statistical models with a large number of parameters. The values for these parameters are determined in a process called training from a large bilingual corpus consisting of parallel English-French text. For the training of the translation model, the English and French text is first aligned on a sentence by sentence basis. The process of aligning sentences is described in Sections 12 and 13.

## 2 Translation Systems

FIG. 4 depicts schematically a preferred embodiment of a batch translation system 401 which translates text without user assistance. The operation of the system involves three basic steps. In the first step 403 a portion of source text is captured from the input source document to be translated. In the second step 404, the portion of source text captured in step 403 is translated into the target language. In the third step 405, the translated target text is displayed, made available to at least one of a listen and viewing operation, or otherwise made available. If at the end of these three steps the source document has not been fully translated, the system returns to step 403 to capture the next untranslated portion of source text.

FIG. 5 depicts schematically a preferred embodiment of a user-aided translation system 502 which assists a user 503 in the translation of text. The system is similar to the batch translation system 401 depicted FIG. 4. Step 505 measures, receives, or otherwise captures a portion of source text with possible input from the user. Step 506 displays the text directly to the user and may also output it to an electric media, made available to at least one of a listening and viewing operation or elsewhere. In the user-aided system 502 the step 404 of the batch system 401 is replaced by step 501 in which source text is translated into the target language with input from the user 503. In the embodiment depicted in FIG. 5, the system 502 is accompanied by a user-interface 504 which includes a text editor or the like, and which makes available to the user translation dictionary,

thesauruses, and other information that may assist the user 503 in selecting translations and specifying constraints. The user can select and further edit the translations produced by the system.

In step 505 in FIG. 5, the user-aided system, portions of source text may be selected by the user. The user might, for example, select a whole document, a sentence, or a single word. The system might then show the users possible translations of the selected portion of text. For example, if the user selected only a single word the system might show a ranked list of possible translations of that word. The ranks being determined by statistical models that would be used to estimate the probabilities that the source word translates in various manners in the source context in which the source word appears.

In step 506 in FIG. 5 the text may be displayed for the user in a variety of fashions. A list of possible translations may be displayed in a window on a screen, for example. Another possibility is that possible target translations of source words may be positioned above or under the source words as they appear in the context of the source text. In this way the source text could be annotated with possible target translations. The user could then read the source text with the help of the possible translations as well as pick and choose various translations.

A schematic representation of a work station environment, through which the interfaces with a user-aided system is shown in FIG. 6.

FIG. 7 depicts in more detail the step 404 of the batch translation system 401 depicted in FIG. 4. The step 404 is expanded into four steps. In the first step 701 the input source text 707 is transduced to one or more intermediate source structures. In the second step 702 a set of one or more hypothesized target structures are generated. This step 702 makes use of a language model 705 which assigns probabilities or scores to target structures and a translation model 706 which assigns probabilities or scores to source structures given target structures. In the third step 703 the highest scoring hypothesis is selected. In the fourth step 704 the hypothesis selected in step 703 is transduced into text in the target language 708.

FIG. 8 depicts in more detail the step 501 of the user-aided translation system 502 depicted in FIG. 5 (the user interface has been omitted for simplification). The step 501 is expanded into four steps. In the first step 801 the user 503 is asked to supply zero or more translation constraints or hints. These constraints are used in subsequent steps of the translation process.

In the second step 802 a set of putative translations of the source text is produced. The word putative is used here, because the system will produce a number of possible translations of a given portion of text. Some of these will be legitimate translations only in particular circumstances or domains, and others may simply be erroneous. In the third step these translations are displayed or otherwise made available to the user. In the fourth step 804, the user 503 is asked either to select one of the putative translations produced in step 803, or to specify further constraints. If the user 503 chooses to specify further constraints the system returns to the first step 801. Otherwise the translation selected by the user 503 is displayed or otherwise made available for viewing or the like.

FIG. 9 depicts in more detail the step 802 of FIG. 8 in which a set of putative translations is generated. The step 802 is expanded into four steps. In the first step 901, the input source text is transduced to one or more intermediate target structures. This step is similar to step 701 of FIG. 7 except that here the transduction takes place in a manner

which is consistent with the constraints 906 which have been provided by the user 503. In the second step 902, a set of one or more hypotheses of target structures is generated. This step is similar to step 702 except that here only hypotheses consistent with the constraints 906 are generated. In the third step 903 a set of the highest scoring hypothesis is selected. In the four step 904 each of the hypotheses selected in step 903 is transduced into text in the target language. The sections of text in the target language which result from this step are the output of this module.

A wide variety of constraints may be specified in step 801 of FIG. 8 and used in the method depicted in FIG. 9. For example, the user 503 may provide a constraint which affects the generating of target hypotheses in step 902 by insisting that a certain source word be translated in a particular fashion. As another example, he may provide a constraint which affects the transduction of source text into an intermediate structure in step 901 by hinting that a particular ambiguous source word should be assigned a particular part of speech. As other examples, the user 503 may a) insist that the word 'transaction' appear in the target text; b) insist that the word 'transaction' not appear in the target text; c) insist that the target text be expressed as a question, or d) insist that the source word 'avoir' be translated differently from the way it was translated in the fifth putative translation produced by the system. In other embodiments which transduce source text into parse trees, the user 503 may insist that a particular sequence of source words by syntactically analyzed as a noun phrase. These are just a few of many similar types of examples. In embodiments in which case frames are used in the intermediate source representation, the user 503 may insist that a particular noun phrase fill the donor slot of the verb 'to give'. These constraints can be obtained from a simple user interface, which someone skilled in the art of developing interfaces between users and computers can easily develop.

Steps 702 of FIG. 7 and 902 of FIG. 9 both hypothesize intermediate target structures from intermediate source structures produced in previous steps.

In the embodiments described here, the intermediate source and target structures are sequences of morphological units, such as roots of verbs, tense markers, prepositions and so forth. In other embodiments these structures may be more complicated. For example, they may comprise parse trees or case-frame representations structures. In a default condition, these structures may be identical to the original texts. In such a default condition, the transductions steps are not necessary and can be omitted.

In the embodiments described here, text in the target language is obtained by transducing intermediate target structures that have been hypothesized in previous steps. In other embodiments, text in the target language can be hypothesized directly from source structures. In these other embodiments steps 704 and 904 are not necessary and can be omitted.

An example of a workstation environment depicting a hardware implementation using the translation system (application) in connection with the present invention is shown in FIG. 10. A computer platform 1014 includes a hardware unit 1004, which includes a central processing unit (CPU) 1006, a random access memory (RAM) 1005, and an input/output interface 1007. The RAM 1005 is also called main memory.

The computer platform 1014 typically includes an operating system 1003. A data storage device 1002 is also called a secondary storage and may include hard disks and/or tape drives and their equivalents. The data storage device 1002 represents non-volatile storage. The data storage device 1002 may be used to store data for the language and translation models components of the translation system 1001.

Various peripheral components may be connected to the computer platform 1014, such as a terminal 1012, a microphone 1008, a keyboard 1013, a scanning device 1009, an external network 1010, and a printing device 1011. A user 503 may interact with the translation system 1001 using the terminal 1012 and the keyboard 1013, or the microphone 1008, for example. As another example, the user 503 might receive a document from the external network 1010, translate it into another language using the translation system 1001, and then send the translation out on the external network 1010.

In a preferred embodiment of the present invention, the computer platform 1014 includes an IBM System RS/6000 Model 550 workstation with at least 128 Megabytes of RAM. The operating system 1003 which runs thereon is IBM AIX (Advanced Interactive Executive) 3.1. Many equivalents to this structure will readily become apparent to those skilled in the art.

The translation system can receive source text in a variety of known manners. The following are only a few examples of how the translation system receives source text (e.g. data), to be translated. Source text to be translated may be directly entered into the computer system via the keyboard 1013. Alternatively, the source text could be scanned in using the scanner 1009. The scanned data could be passed through a character recognizer in a known manner and then passed on to the translation system for translation. Alternatively, the user 503 could identify the location of the source text in main or secondary storage, or perhaps on removable secondary storage (such as on a floppy disk), the computer system could retrieve and then translate the text accordingly. As a final example, with the addition of a speech recognition component, it would also be possible to speak into the microphone 1008, have the speech converted into source text by the speech recognition component.

Translated target text produced by the translation application running on the computer system may be output by the system in a variety of known manners. For example, it may be displayed on the terminal 1012, stored in RAM 1005, stored data storage 1002, printed on the printer 1011, or perhaps transmitted out over the external network 1010. With the addition of a speech synthesizer it would also be possible to convert translated target text into speech in target language.

Step 403 in FIG. 4 and step 505 in FIG. 5 measure, receive or otherwise capture a portion of source text to be translated. In the context of this invention, text is used to refer to sequences of characters, formatting codes, the typographical marks. It can be provided to the system in a number of different fashions, such as on a magnetic disk, via a computer network, as the output of an optical scanner, or of a speech recognition system. In some preferred embodiments, the source text is captured a sentence at a time. Source text is parsed into sentences using a finite-state machine which examines the text for such things as upper and lower case characters and sentence terminal punctuations. Such a machine can easily be constructed by someone skilled in the art. In other embodiments, text may be parsed into units such as phrases or paragraphs which are either smaller or larger than individual sentences.

### 3 Source Transducers

In this section, some embodiments of the source-transducer 701 will be explained. The role of this transducer is to produce one or more intermediate source-structure representations from a portion of text in the source language.

### 3.1 Overview

An embodiment of the source-transducer 701 is shown in FIG. 11. In this embodiment, the transducer takes as input a sentence in the source language and produces a single intermediate source-structure consisting of a sequence of linguistic morphs. This embodiment of the source-transducer 701 comprises transducers that:

- tokenize raw text 1101;
- determine words from tokens 1102;
- annotate words with parts of speech 1103;
- perform syntactic transformations 1104;
- perform morphological transformations 1105;
- annotate morphs with sense labels 1106.

It should be understood that FIG. 11 represents only one embodiment of the source-transducer 701. Many variations are possible. For example, the transducers 1101, 1102, 1103, 1104, 1105, 1106 may be augmented and/or replaced by other transducers. Other embodiments of the source-transducer 701 may include a transducer that groups words into compound words or identifies idioms. In other embodiments, rather than a single intermediate source-structure being produced for each source sentence, a set of several intermediate source-structures together with probabilities or scores may be produced. In such embodiments the transducers depicted in FIG. 11 can be replaced by transducers which produce at each stage intermediate structures with probabilities or scores. In addition, the intermediate source-structures produced may be different. For example, the intermediate structures may be entire parse trees, or case frames for the sentence, rather than a sequence of morphological units. In these cases, there may be more than one intermediate source-structure for each sentence with scores, or there may be only a single intermediate source-structure.

#### 3.2 Components

Referring still to FIG. 11, the transducers comprising the source-transducer 701 will be explained. For concreteness, these transducers will be discussed in cases in which the source language is either English or French.

##### 3.2.1 Tokenizing Transducers

The purpose of the first transducer 1101, which tokenizes raw text, is well illustrated by the following Socratic dialogue:

How do you find words in text?  
 Words occur between spaces.  
 What about "however,"? Is that one word or two?  
 Oh well, you have to separate out the commas.  
 Periods too?  
 Of course.  
 What about "Mr."?  
 Certain abbreviations have to be handled separately.  
 How about "shouldn't"? One word or two?  
 One.  
 So "shouldn't" is different from "should not"?  
 Yes.  
 And "Gauss-Bonnet", as in the "Gauss-Bonnet Theorem"?  
 Two names, two words.  
 So if you split words at hyphens, what do you do with "vis-à-vis"?  
 One word, and don't ask me why.  
 How about "stingray"?  
 One word, of course.  
 And "manta ray"?  
 One word: it's just like stingray.  
 But there's a space.  
 Too bad.  
 How about "inasmuch"?

Two.

Are you sure?

No.

This dialogue illustrates that there is no canonical way of breaking a sequence of characters into tokens. The purpose of the transducer 1101, that tokenizes raw text, is to make some choice.

In an embodiment in which the source-language is English, this tokenizing transducer 1101 uses a table of a few thousand special character sequences which are parsed as single tokens and otherwise treat spaces and hyphens as word separators. Punctuation marks and digits are separated off and treated as individual words. For example, 87, is tokenized as the three words 8 7 and ,.

In another embodiment in which the source-language is French, the tokenizing transducer 1101 operates in a similar way. In addition to space and hyphen, the symbol -t- is treated as a separator when tokenizing French text.

##### 3.2.2 Token-Word Transducers

The next transducer 1102 determines a sequence of words from a sequence of token spellings. In some embodiments, this transducer comprises two transducers as depicted in FIG. 12. The first of these transducers 1201 determines an underlying case pattern for each token in the sequence. The second of these transducers 1202 performs a number of specialized transformations. These embodiments can easily be modified or enhanced by a person skilled in the art.

##### 3.2.3 True-Case Transducers

The purpose of a true-case transducer 1201 is made apparent by another Socratic dialogue:

When do two sequences of characters represent the same word?  
 When they are the same sequences.  
 So, "the" and "The" are different words?  
 Don't be ridiculous. You have to ignore differences in case.  
 So "Bill" and "bill" are the same word?  
 No. "Bill" is a name and "bill" is something you pay.  
 With proper names the case matters.  
 What about the two "May"'s in "May I pay in May?"  
 The first one is not a proper name. It is capitalized because it is the first word in the sentence.  
 Then, how do you know when to ignore case and when not to?  
 If you are human, you just know.

Computers don't know when case matters and when it doesn't. Instead, this determination can be performed by a true-case transducer 1201. The input to this transducer is a sequence of tokens labeled by a case pattern that specified the case of each letter of the token as it appears in printed text. These case patterns are corrupted versions of true-case patterns that specify what the casing would be in the absence of typographical errors and arbitrary conventions (e.g., capitalization at the beginning of sentences). The task of the true-case transducer is to uncover the true-case patterns.

In some embodiments of the true-case transducer the case and true-case patterns are restricted to eight possibilities

---

L\* U\* UL\* ULUL\* ULLUL\* UUL\* UUUL\* LUL\*

---

Here U denotes an upper case letter, L a lower case letter, U\* a sequence of one or more upper case letters, and L\* a sequence of one or more lower case letters. In these embodiments, the true-case transducer can determine a true-case of an occurrence of a token in text using a simple method comprising the following steps:

## 15

1. Decide whether the token is part of a name. If so, set the true-case equal to the most probable true-case beginning with a U for that token.
2. If the token is not part of a name, then check if the token is a member of a set of tokens which have only one true-case. If so, set the true-case appropriately.
3. If the true-case has not been determined by steps 1 or 2, and the token is the first token in a sentence then set the true-case equal to the most probable true-case for that token.
4. Otherwise, set the true-case equal to the case for that token.

In an embodiment of the true-case transducer used for both English and French, names are recognized with a simple finite-state machine. This machine employs a list of 12,937 distinct common last names and 3,717 distinct common first names constructed from a list of 125 million full names obtained from the IBM online phone directory and a list of names purchased from a marketing corporation. It also uses lists of common precursors to names (such as Mr., Mrs., Dr., Mlle., etc.) and common followers to names (such as Jr., Sr., III, etc.).

The set of tokens with only one true-case consists of all tokens with a case-pattern entropy of less than 1.25 bits, together with 9,506 Number of records in (o+p)8260.lcwtab English and 3,794 records in r8260.coerce(a+b+c) French tokens selected by hand from the tokens of a large bilingual corpus. The most probable case pattern for each English token was determined by counting token-case cooccurrences in a 67 million word English corpus, and for each French token by counting cooccurrences in a 72 million word French corpus. (In this counting occurrences at the beginnings of sentences are ignored).

### 3.2.4 Specialized Transformation Transducers

Referring still to FIG. 12, the transducer 1202 performs a few specialized transformations designed to systematize the tokenizing process. These transformations can include, but are not limited to, the correction of typographical errors, the expansion of contractions, the systematic treatment of possessive, etc. In one embodiment of this transducer for English, contractions such as don't are expanded to do not, and possessives such as John's and nurses' are expanded to John's and nurses'. In one embodiment of this transducer for French, sequences such as s'il, qu'avez, and j'adore are converted to pairs of tokens such as si il, que avez, and je adore. In addition, a few thousand sequences such as afin de are contracted to strings such as afin\_de. These sequences are obtained from a list compiled by a native French speaker who felt that such sequences should be treated as individual words. Also the four strings au, aux, du, and des are expanded to à le, à les, de le, and de les respectively.

### 3.2.5 Part-of-Speech Labeling Transducers

Referring again to FIG. 11, the transducer 1103 annotates words with part-of-speech labels. These labels are used by the subsequent transducers depicted the figure. In some embodiments of transducer 1103, part-of-speech labels are assigned to a word sequence using a technique based on hidden Markov models. A word sequence is assigned the most probable part-of-speech sequence according to a statistical model, the parameters of which are estimated from large annotated texts and other even larger un-annotated texts. The technique is fully explained in article by Bernard Merialdo entitled 'Tagging text with a Probabilistic Model' in the *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*, May 14-17, 1991. This article is incorporated by reference herein.

## 16

In an embodiment of the transducer 1103 for tagging of English, a tag set consisting of 163 parts of speech is used. A rough categorization of these parts of speech is given in Table 1.

In an embodiment of this transducer 1103 for the tagging of French, a tag set consisting of 157 parts of speech is used. A rough categorization of these parts of speech is given in Table 2.

TABLE 1

Parts of Speech for English	
29	Nouns
27	Verbs
20	Pronouns
17	Determiners
16	Adverbs
12	Punctuation
10	Conjunctions
8	Adjectives
4	Prepositions
20	Other

TABLE 2

Parts-of-Speech for French	
105	Pronouns
26	Verbs
18	Auxiliaries
12	Determiners
7	Nouns
4	Adjectives
4	Adverbs
4	Conjunctions
2	Prepositions
2	Punctuation
12	Other

### 3.2.6 Syntactic Transformation Transducers

Referring still to FIG. 11, the transducer 1104, which performs syntactic transformations, will be described.

One function of this transducer is to simplify the subsequent morphological transformations 1105. For example, in morphological transformations, verbs may be analyzed into a morphological unit designating accidence followed by another unit designating the root of the verb. Thus in French, the verb ira might be replaced by 3s\_future\_indicative aller, indicating that ira is the third person singular of the future tense of aller. The same kind of transformation can be performed in English by replacing the sequence will go by future\_indicative to\_go. Unfortunately, often the two words in such a sequence are separated by intervening words as in the sentences:

will be go play in the traffic?

he will not go to church.

Similarly, in French the third person of the English verb went is expressed by the two words est allé, and these two words can be separated by intervening words as in the sentences:

est-t-il allé?

Il n'est pas allé.

It is possible to analyze such verbs morphologically with simple string replacement rules if various syntactic transformations that move away intervening words are performed first.

A second function of the syntactic transducer 1104 is to make the task presented to the statistical models which generate intermediate target-structures for an intermediate source-structure as easy as possible. This is done by performing transformations that make the forms of these structures more similar. For example, suppose the source lan-

guage is English and the target language is French. English adjectives typically precede the nouns they modify whereas French adjectives typically follow them. To remove this difference, the syntactic transducer 1104 includes a transducer which moves French words labeled as adjectives to positions proceeding the nouns which they modify.

These transducers only deal with the most rudimentary linguistic phenomena. Inadequacies and systematic problems with the transformations are overcome by the statistical models used later in the target hypothesis-generating module 702 of the invention. It should be understood that in other embodiments of the invention more sophisticated schemes for syntactic transformations with different functions can be used.

The syntactic transformations performed by the transducer 1104 are performed in a series of steps. A sequence of words which has been annotated with parts of speech is fed into the first transducer. This transducer outputs a sequence of words and a sequence of parts of speech which together serve as input to the second transducer in the series, and so forth. The word and part-of-speech sequences produced by the final transducer are the input to the morphological analysis transducer.

### 3.2.7 Syntactic Transducers for English

FIG. 13, depicts an embodiment of a syntactic transducer 1104 for English. Although in much of this document, examples are described in which the source language is French and the target language is English, here for reasons of exposition, an example of a source transducer for a source language of English is provided. In the next subsection, another example in which the source language is French is provided. Those with a basic knowledge of linguistics for other language will be able to construct similar syntactic transducer for those languages.

The syntactic transducer in FIG. 13 is comprised of three transducer that

- perform question inversion 1301;
- perform do-not coalescence 1302;
- perform adverb movement 1303.

To understand the function of the transducer 1301 that performs question inversion, note that in English questions the first auxiliary verb is often separated by a noun phrase from the root of the verb as in the sentences:

- does that elephant eat?
- which car is he driving?

The transducer 1301 inverts the auxiliary verb with the subject of the sentence, and then converts the question mark to a special QINV marker to signal that this inversion has occurred. For example, the two sentences above are converted by this transducer to:

- that elephant eats QINV
- which car he is driving QINV

This transducer also removes supporting do's as illustrated in the first sentence.

To understand the function of the transducer 1302 that performs do-not coalescence, note that in English, negation requires the presence of an auxiliary verb. When one doesn't exist, an inflection of to do is used. The transducer 1302 coalesces the form of to do with not into the string do\_not. For example,

---

John does not like turnips.  
⇒ John do\_not like turnips.

---

The part of speech assigned to the main verb, like above, is modified by this transducer to record the tense and person of

to do in the input. When adverbs intervene in emphatic sentences the do\_not is positioned after the intervening adverbs:

---

John does really not like turnips.  
⇒ John really do\_not like turnips.

---

To understand the function of the final transducer 1303 that performs adverb movement, note that in English, adverbs often intervene between a verb and its auxiliaries. The transducer 1303 moves adverbs which intervene between a verb and its auxiliaries to positions following the verb. The transducer appends a number onto the adverbs it moves to record the positions from which they are moved. For example,

---

Iraq will probably not be completely balkanized.  
⇒ Iraq will\_be\_en to\_balkanize probably\_M2 not\_M2 completely\_M3.

---

An M2 is appended to both probably and to not to indicate that they originally preceded the second word in the verbal sequence will be balkanized. Similarly, an M3 is appended to completely to indicate that it preceded the third word in the verbal sequence.

The transducer 1303 also moves adverbs that precede verbal sequences to positions following those sequences:

---

John really eats like a hog.  
⇒ John eats really\_M1 like a hog.

---

This is done in order to place verbs close to their subjects.

### 3.2.8 Syntactic Transducers for French

Referring now to FIG. 14, an embodiment of the syntactic transducer 1104 for French is described. This embodiment comprises four transducers that

- perform question inversion 1401;
- perform discontinuous negative coalescence 1402;
- perform pronoun movement 1403;
- perform adverb and adjective movement 1104.

The question inversion transducer 1401 undoes French question inversion much the same way that the transducer 1301 undoes English question inversion:

---

mangez-vous des légumes?  
⇒ vous mangez des légumes QINV  
où habite-t-il?  
⇒ où il habite QINV  
le lui avez-vous donné?  
⇒ vous le lui avez donnéQINV

---

This transducer 1401 also modifies French est-ce que questions:

---

est-ce qu'il mange comme un cochon?  
⇒ il mange comme un cochon EST-CE\_QUE

---



To understand the function of the transducer 1402, which performs negative coalescence, note that propositions in French are typically negated or restricted by a pair of words, ne and some other word, that surround the first word in a verbal sequence. The transducer 1402 moves the ne next to its mate, and then coalesces the two into a new word:

---

```

je ne sais pas.
⇒ je sais ne_pas.
Jean n'a jamais mangé comme un cochon.
⇒ Jean a ne_jamais mangé comme un cochon.
il n'y en a plus.
⇒ il y en a ne_plus.

```

---

To understand the function of the transducer 1403, which performs pronoun movement, note that in French, direct-object, indirect-object and reflexive pronouns usually precede verbal sequences. The transducer 1403 moves these pronouns to positions following these sequences. It also maps these pronouns to new words that reflect their roles as direct-object or indirect-objects, or reflexive pronouns. So, for example, in the following sentence le is converted to le DPRO because it functions as a direct object and vous to vous IPRO because it functions as an indirect object:

---

```

je vous le donnerai.
⇒ je donnerai le_DPRO vous_IPRO.

```

---

In the next sentence, vous is tagged as a reflexive pronoun and therefore converted to

---

```

vous_RPRO.      vous vous lavez les mains.
⇒ vous lavez vous_RPRO les mains.

```

---

The allative pronominal clitic y and the ablative pronominal clitic en are mapped to the two-word tuples à y\_PRO and de en\_PRO:

---

```

je y penserai.
⇒ je penserai à y_PRO.
j'en ai plus.
⇒ je ai plus de en_PRO.

```

---

The final transducer 1404 moves adverbs to positions following the verbal sequences in which they occur. It also moves adjectives to positions preceding the nouns they modify. This is a useful step in embodiments of the present invention that translate from French to English, since adjectives typically precede the noun in English.

### 3.2.9 Morphological Transducer

Referring again to FIG. 11, a transducer 1105, which performs morphological transformations, will be described. One purpose of this transducer is to make manifest in the intermediate source-structure representation the fraternity of different forms of a word. This is useful because it allows for more accurate statistical models of the translation process. For example, a system that translates from French to English but does not use a morphological transducer can not benefit from the fact that sentences in which parle is translated as speaks provide evidence that parlé should be translated as spoken. As a result, parameter estimates for rare words are inaccurate even when estimated from a very large training

sample. For example, even in a sample from the Canadian Parlement of nearly 30 million words of French text, only 24 of the 35 different spellings of single-word inflections of the verb parler actually occurred.

A morphological transducer 1104 is designed to ameliorate such problems. The output of this transducer is a sequence of lexical morphemes. These lexical morphemes will sometimes be referred to in this application as morphological units or simply morphs. In an embodiment of transducer 1104 used for English, inflection morphological transformations are performed that make evident common origins of difference conjugations of the same verb; the singular and plural forms of the same noun; and the comparative and superlative forms of adjectives and adverbs. In an embodiment of transducer 1104 used for French, morphological inflectional transformations are performed that make manifest the relationship between conjugations of the same verb; and forms of the same noun or adjective different in gender and number are performed. These morphological transformations are reflected in the sequence of lexical morphemes produced. The examples below illustrate the level of detail in these embodiments of a morphological transducer 1104:

---

```

he was eating the peas more quickly than I.
⇒ he V_past_progressive to_eat the pea N_PLURAL
quick er_ADV than I.
nous en mangeons rarement.
⇒ nous V_present_indicative_1p manger rare ment_ADV
de en_PRO
ils se sont lavés les mains sales.
⇒ ils V_past_3p laver se_RPRO les sale main N_PLURAL.

```

---

### 3.2.10 Sense-Labeling Transducers

Referring again to FIG. 11, the transducer 1106, which annotates a lexical morph sequence produced by the transducer 1105 with part-of-speech labels, will be explained. Much of the allure of the statistical approach to transfer in machine translation is the ability of that approach to formally cope with the problem of lexical ambiguity. Unfortunately, statistical methods are only able to mount a successful attack on this problem when the key to disambiguating the translation of a word falls within the local purview of the models used in transfer.

Consider, for example, the French work prendre. Although prendre is most commonly translated as to take, it has a number of other less common translations. A trigram model of English can be used to translate Je vais prendre la décision as I will make the decision because the trigram make the decision is much more common than the trigram take the decision. However, a trigram model will not be of much use in translating Je vais prendre ma propre décision as I will make my own decision because in this case take and decision no longer fall within a single trigram.

In the paper, "Word Sense Disambiguation using Statistical Methods" in the proceedings of the 29th Annual Meeting of the Association for Computational Linguistics, published in June of 1991 by the Association of Computational Linguistics and incorporated by reference herein, a description is provided of a method of asking a question about the context in which a word appears to assign that word a sense. The question is constructed to have high mutual information with the translation of that word in its context. By modifying the lexical entries that appear in a sequence of morphemes to reflect the senses assigned to these entries, informative global information can be encoded locally and thereby made available to the statistical models used in transfer.



Although the method described in the aforementioned paper assigns senses to words, the same method applies equally well to the problem of assigning senses to morphemes, and is used here in that fashion. This transducer 1106, for example maps prendre to prendre 1 in the sentence

Je vais prendre ma propre voiture. ps but to prendre 2 in the sentence

Je vais prendre ma propre décision.

It should be understood that other embodiments of the sense-labelling transducer are possible. For example, the sense-labelling can be performed by asking not just a single question about the context but a sequence of questions arranged in a decision tree.

### 3.3 SOURCE-TRANSDUCERS WITH CONSTRAINTS

In some embodiments, such as that depicted in FIG. 9, a source-structure transducer, such as that labelled 901, accepts a set of constraints that restricts its transformations source text to an intermediate target structure in source text.

Such constraints include, but are not limited to, requiring that a particular phrase be translated as a certain linguistic component of a sentence, such a noun-phase; requiring that a source word be labelled as a certain part-of-speech such as a verb or determiner; requiring that a source word be morphologically analyzed a certain way; requiring that a source word be annotated with a particular sense label;

in embodiments in which the intermediate structure encodes parse-tree or case-frame information, requiring a certain parse or case-frame structure for a sentence; morphologically analyzed in a particular way, or be annotated with a particular sense-label,

A source-transducer accepting such constraints in similar to source transducers as described in this section. Based on the descriptions already given, such transducers can be constructed by a person with a computer science background and skilled in the art.

## 4 FINITE-STATE TRANSDUCERS

This section provides a description of an embodiment of a mechanism by which the syntactic transductions in step 1104 and the morphological transductions in step 1105 are performed. The mechanism is described in the context of a particular example depicted in 15.. One with a background in computer science and skilled in the art of producing finite-state transducers can understand from this example how to construct syntactic and morphological transducers of the type described above.

The example transducer inverts questions involving do, does, and did. After steps 1101, 1102, and 1103, the source text Why don't you ever succeed? is transduced into parallel word and part-of-speech sequences 1501:

why	do	not	you	ever	succeed	?
RRQ	VD0	XX	PPY	RR	VV0	?

Here, RRQ and RR are adverb tags, VD0 and VV0 are verb tags, XX is a special tag for the word not, PPY is a pronoun tag, and ? is a special tag for a question mark.

The example transducer converts these two parallel sequences to the parallel word and part-of-speech sequences 1502:

why	you	succeed	do_not_M0	ever_M1	QINV
RRQ	PPY	VV0	XX	RR	QINV

Here, QINV is a marker which records the fact that the original input sentence was question inverted.

A mechanism by which a transducer achieves this transformation is depicted in FIG. 16, and is comprised of four components:

- a finite-state pattern matcher 1601;
- an action processing module 1603;
- a table of patterns 1602;
- a table of action specifications 1503.

The transducer operates in the following way:

1. One or more parallel input sequences 1605 are captured by the finite-state pattern-matcher 1601;
2. The finite-state pattern-matcher compares the input sequences against a table of patterns 1602 of input sequences stored in memory;
3. A particular pattern is identified, and an associated action-code 1606 is transmitted to the action-processing module 1603;
4. The action-processing module contains a specification of the transformation associated to this action code from a table of actions 1503 stored in memory;
5. The action-processing module applies the transformation to the parallel input streams to produce one or more parallel output sequences 1604;

The parallel input streams captured by the finite-state pattern matcher 1601 are arranged in a sequence of attribute tuples. An example of such a sequence is the input sequence 1501 depicted in FIG. 15. This sequence consists of a sequence of positions together with a set of one or more attributes which take values at the positions. A few examples of such attributes are the token attribute, the word attribute, the case-label attribute, the part-of-speech attribute, the sense-label attribute. The array of attributes for a given position will be called an attribute tuple.

For example, the input attribute tuple sequence 1501 in FIG. 15 is seven positions long and is made up of two dimensional attribute tuples. The first component of an attribute tuple at a given position refers to the word attribute. This attribute specifies the spellings of the words at given positions. For example, the first word in the sequence 1501 is why. The second component of an attribute tuple at a given position for this input sequence refers to a part-of-speech tag for that position. For example, the part of speech at the first position is RRQ. The attribute tuple at position 1 is thus the ordered pair why, RRQ.

The parallel output streams produced by the action processing module 1603 are also arranged as a sequence of attribute tuples. The number of positions in an output sequence may be different from the number of positions in an input sequence.

For example, the output sequence 1502 in FIG. 15, consists of six positions. Associated with each position is a two-dimensional attribute tuple, the first coordinate of which is a word attribute and the second coordinate of which is a part-of-speech attribute.

An example of a table of patterns 1602 is shown in FIG. 17. This table is logically divided into a number of parts or blocks.

Pattern-Action Blocks. The basic definitions of the matches to be made and the actions to be taken are contained in pattern-action blocks. A pattern-action block comprises of a list of patterns together with the name of actions to be invoked when patterns in an input attribute-tuple sequence 1605 are matched.

Auxiliary Pattern Blocks. Patterns that can be used as sub-patterns in the patterns of pattern-action blocks are defined in Auxiliary Pattern blocks. Such blocks contain lists of labelled patterns of attributes tuples. These labelled patterns do not have associated actions, but can be referenced by their name in the definitions of other patterns.

In FIG. 17 there is one Auxiliary Pattern block. This block defines four auxiliary patterns. The first of these has a name ADVERB and matches single tuple adverb-type constructions. The second has a name of BARE NP and matches certain noun-phrase-type constructions. Notice that this auxiliary pattern makes use of the ADVERB pattern in its definition. The third and fourth auxiliary patterns match other types of noun phrases.

Set Blocks. Primary and auxiliary patterns allow for sets of attributes. In FIG. 17, for example, there is a set called DO SET, of various forms of to do, and another set PROPER NOUN TAG of proper-noun tags.

Patterns are defined in terms of regular expressions of attribute tuples. Any pattern of attribute tuples that can be recognized by a deterministic finite-state automata can be specified by a regular expression. The language of regular expressions and methods for constructing finite-state automata are well known to those skilled in computer science. A method for constructing a finite-state pattern matcher from a set of regular expressions is described in the article "LEX—A lexical Analyzer Generator," written by Michael E. Lesk, and appearing in the Bell Systems Technical Journal, Computer Science Technical Report Number. 39, published in October of 1975.

Regular expressions accepted by the pattern matcher 1601 are described below.

Regular Expressions of Attribute Tuples: A regular expression of attribute tuples is a sequence whose elements are either

1. an attribute tuple;
2. the name of an auxiliary regular expression; or
3. a register name.

These elements can be combined using one of the logical operations:

Operator	Meaning	Usage	Matches
.	concatenation	A.B	A followed by B
	union (i.e. or)	A B	A or B
*	0 or more	A*	0 or more A's
?	0 or 1	A?	0 or 1 A's
+	1 or more	A+	1 or more A's

Here A and B denote other regular expressions. Examples of these logical operations are:

Expression	Matches
A?.B.C	0 or 1 A's followed by B then by C
(A*)(B+)	0 or more A's or 1 or more B's

-continued

Expression	Matches
(A B).C	A or B, followed by C

Attribute Tuples: The most common type of element in a regular expression is an attribute tuple. An attribute tuple is a vector whose components are either:

1. an attribute (as identified by its spelling);
2. a name of a set of attributes;
3. the wild card attribute.

These elements are combined using the following operators:

Operator	Meaning	Usage
.	Delimiter between coordinates of an attribute tuple	a,b
-	Negation	~a
#	Wild Card	#

(Here a and b denote attribute spellings or attribute set names).

The meanings of these operators are best illustrated by example. Let a, b, and c denote either attribute spellings or set names. Assume the dimension of the attribute tuples is 3. Then:

Attribute Tuple	Matches
a,b,c	First attribute matches a, second match b, third matches c
,b,c	First attribute elided (matches anything), Second attribute matches b, third matches c
,b,	First and third attribute elided (match anything) Second attribute matches b
a	Second and third attributes elided (Match anything) First matches a
#,b,	First attribute wild-card (i.e. matches anything) Second attribute matches b. Third attribute elided
a,~,c	Second attribute matches anything EXCEPT b. Third matches anything EXCEPT c.

Auxiliary Regular Expressions: A second type of element in a regular expression is an auxiliary regular expression. An auxiliary regular expression is a labelled regular expression which is used as a component of a larger regular expression.

Logically, a regular expression involving auxiliary regular expressions is equivalent to the regular expression obtained by resolving the reference to the auxiliary pattern. For example, suppose an auxiliary regular expression named D has been defined by:

$$D=A.B+A^*$$

where A,B denote attribute tuples (or other auxiliary patterns). Then:

Expression	is equivalent to
C.D	C.A.B+A*

-continued

Expression	is equivalent to
D+.C.D	(A.B+.A*)+.A.B+.A*.C.A*.B+.A*

Registers: Just knowing that a regular expression matches an input attribute tuple sequence usually does not provide enough information for the construction of an appropriate output attribute tuple sequence. Data is usually also required about the attribute tuples matched by different elements of the regular expression. In ordinary LEX, to extract this type of information often requires the matched input sequence to be parsed again. To avoid this cumbersome approach, the pattern-matcher 1601 makes details about the positions in the input stream of the matched elements of the regular expression more directly available. From these positions, the identities of the attribute tuples can then be determined.

Positional information is made available through the use of registers. A register in a regular expression does not match any input. Rather,

1. After a match, a register is set equal to the position in the input sequence of the next tuple in the input sequence that is matched by an element of the regular expression to the right of the register.
2. If no further part of the regular expression to the right of the register matches, then the register is set equal to zero,

The operation of registers is best illustrated by some examples. These examples use registers [1] and [2]:

Expression	Contents of Registers after match
A.[1].B.C	Reg 1: First position of B match
A.[2].(C D)	Reg 2: First position of either C or D match
A.[1].B*.[2].C	Reg 1: If B matches: First position of B match Otherwise: First position of C match
A.[1].B*.C*	Reg 2: First position of C match Reg 1: If B matches: First position of B match If C matches: First position of C match Otherwise: 0 Reg 2: If C matches: First position of C match Otherwise: 0

A pattern-action block defines a pattern matcher. When an input attribute-tuple sequence is presented to the finite-state pattern matcher a current input position counter is initialized to 1 denoting that the current input position is the first position of the sequence. A match at the current input position is attempted for each pattern. If no pattern matches, an error occurs. If more than one pattern matches, the match of the longest length is selected. If several patterns match of the same longest length, the one appearing first in the definition of the pattern-action block is selected. The action code associated with that pattern is then transmitted to the action processing module 1603.

Transformations by the action processing module are defined in a table of actions 1503 which is stored in memory. The actions can be specified in specified in any one of a number of programming languages such as C, PASCAL, FORTRAN, or the like.

In the question-inversion example, the action specified in the pseudo-code in FIG. 18 is invoked when the pattern defined by the regular expression in lines 3-4 is matched. This action inverts the order of the words in certain ques-

tions involving forms of do. An instance of this action is shown in FIG. 15. In the pseudo-code of FIG. 18 for this action, the symbol @reg(i) denotes the contents of register i. In line 6 of this pseudo-code, the output attribute tuple sequence is set to null.

A question matched by the regular expression in lines 3-4 may or may not begin with a (so-called) wh- word in the set WH NP. If it does match, the appropriate action is to append the input tuple in the first position to the output sequence. This is done in lines 8-9.

After the wh- word, the next words of the output sequence should be the subject noun phrase of the input sequence. This is made so in line 11-12 that appends all tuples matching the regular expression SUBJECT NP to the output sequence.

For negative questions involving forms of do, the part-of-speech tag of the output verb and of the output adverbs are the same as those of the input verb and adverbs. Thus the entire input tuple sequences corresponding to these words can be appended to the output. This is done in lines 15-18.

For positive questions the tag attribute of the output verb may be different than that of the input verb. This is handled in lines 25-37. The input word attribute for the verb is appended to the output word attribute in lines 26 and 31 and 35. The output tag attribute is selected based on the form of do in the input sequence. Explicit tag values are appended to the output sequence in lines 32 and 37.

The remaining input words and tags other than the question mark are written to the output sequence in lines 43-44. The input sequence is completed in line 46 by the marker QINV signalling question inversion, together with the appropriate tag.

## 5 TARGET TRANSDUCERS

In this section, some embodiments of the the target-structure transducer 704 will be explained. The role of this transducer is to produce portions of text in the target language from one or more intermediate target-structures.

### 5.1 OVERVIEW

An embodiment of a target-structure transducer 704 for English is shown in FIG. 19. In this embodiment, the transducer 704 converts an intermediate target-structure for English consisting of a sequence of linguistic morphs into a single English sentence. This embodiment performs the inverse transformations of those performed by the source transducer 702 depicted in FIG. 13. That is, if the transducer 702 is applied to an English sentence, and then the transducer 704 is applied to the sequence of linguistic morphs produced by the transducer 702 the original English sentence is recovered.

This embodiment of the target-transducer 704 comprises a sequence of transducers which:

- generate nouns, adverbs and adjectives from morphological units 1901;
- provide do-support for negatives 1902;
- separate compound verb tenses 1903;
- provide do-support for inverted questions 1904;
- conjugate verbs 1905;
- reposition adverbs and adjectives 1906;
- restore question inversion 1907;
- resolve ambiguities in the choice of target words resulting from ambiguities in the morphology 1908; and
- case the target sentence 1909.

It should be understood that FIG. 19 represents only one possible embodiment of the target-transducer 704. Many variations are possible. For example, in other embodiments, rather than a single target sentence being produced for each intermediate target-structure, a set of several target sentences together with probabilities or scores may be produced. In such embodiments, the transducers depicted in FIG. 19 can be replaced by transducers which produce at each stage several target sentences with probabilities or scores. Moreover, in embodiments of the present invention in which the intermediate target-structures are more sophisticated than lexical morph sequences, the target-structure transducer is also more involved. For example, if the intermediate target-structure consists of a parse tree of a sentence or case frames for a sentence, then the target-structure transducer converts these to the target language.

### 5.2 COMPONENTS

The transducers depicted in FIG. 19 will now be explained. The first seven of these transducers 1901 1902 1903 1904 1905 1906 1907 and the last transducer 1909 are implemented using a mechanism similar to that depicted in FIG. 12 and described in subsection Finite-State Transducers. This is the same mechanism that is used by the syntactic transducers 1104 and the morphological transducers 1105 that are part of the source-transducer 701. Such transducers can be constructed by a person with a background in computer science and skilled in the art. The construction of the transducer 1908 will be explained below.

The first transducer 1901 begins the process of combining lexical morphs into words. It generates nouns, adverbs and adjectives from their morphological constituents. For example, this transducer performs the conversions:

---

```

the ADJ_er big boy N_plural V_bare to_eat rapid ADV_ly.
⇒ the bigger boys V_bare to_eat rapidly.
he V_present_3s to_eat the pea N_PLURAL quick ADV_er than I.
⇒ he V_present_3s to_eat the peas more quickly than I.

```

---

(The final sentences here are The bigger boys eat rapidly and He eats the peas more quickly than I.)

There are sometimes ambiguities in generating adverbs and adjectives from their morphological units. For example, the lexical morph sequence quick ADV ly can be converted to either quick or quickly, as for example in either how quick can you fun? or how quickly can you run?. This ambiguity is encoded by the transducer 1901 by combining the different possibilities into a single unit. Thus, for example, the transducer 1901 performs the conversion:

---

```

how can quick ADV_ly M1 you V_bare to_run ?
⇒ how can quick_quickly M1 you V_bare to_run?

```

---

The second transducer 1902 provides do-support for negatives in English by undoing the do-not coalescences in the lexical morph sequence. Thus, for example it converts:

---

```

John V_present_3s to_eat do_not M1 really M1 a lot.
⇒ John V_present_3s to_do eat not M2 really M2 a lot.

```

---

by splitting do from not and inserting to do at the head of the verb phrase. (The final sentence here is John does not really eat a lot.) The move-markers M1, M2, etc, measure

the position of the adverb (or adjective) from the head of the verb phrase. These markers are correctly updated by the transformation.

The next transducer 1903 provides do-support for inverted questions by inserting the appropriate missing form of the to do into lexical morph sequences corresponding to inverted questions. This is typically inserted between the verb tense marker and the verb. Thus, for example, if converts:

---

```

John V_present_3s to_like to_eat QINV
⇒ John V_present_3s to_do like to_eat QINV

```

---

The next transducer 1904 prepares the stage for verb-conjugation by breaking up complicated verb tenses into simpler pieces.

---

```

John V_has_been_ing to_eat a lot recently
⇒ John V_header has been V_ing to_eat a lot recently.

```

---

This transducer also inserts a verb phrase header marker V header for use by later transducers. In particular the verb phrase header marker is used by the transducer 1907 that repositions adjective and adverb.

The next transducer 1905 conjugates verbs by combining the root of the verb and the tense of the verb into a single word. This process is facilitated by the fact that complicated verb tenses have already been broken up into simpler pieces. For example, the previous sentence is transduced to:

---

```

⇒ John V_header has been eating a lot recently.

```

---

The next transducer 1906 repositions adverbs and adjectives to their appropriate places in the target sentence. This is done using the move-markers M1, M2, etc., that are attached to adjectives. Thus for example, this transducer converts:

---

```

Iraq V_header will be Balkanized probably M2 not M2.
⇒ Iraq V_header will probably not be Balkanized

```

---

The verb phrase header marker V header is kept for use by the next transducer 1907.

The transducer 1907 restores the original order of the noun and verb in inverted questions.

These questions are identified by the question inversion marker QINV. In most cases, to restore the inverted question order, the modal or verb at the head of the verb phrase is moved to the beginning of the sentence:

---

```

John V_header was here yesterday QINV
⇒ was John here yesterday?

```

---

If the sentence begins with one of the words what, who, where, why, whom, how, when, which the modal or verb is moved to the second position in the sentence;

---

where John V\_header was yesterday QINV  
 ⇒ where was John yesterday?

---

A special provision is made for moving the word not in inverted questions. The need to do this is signaled by a special marker M0 in the lexical morph sentence. The not preceding the M0 marker is moved to the position following the modal or verb;

---

he V\_header can swim not M0 QINV.  
 ⇒ can not he swim?  
 John V\_header does eat not M0 potatoes?  
 ⇒ Does not John eat potatoes?

---

After the inverted question order has been restored, the question inversion marker QINV is removed and replaced by a question mark at the end of the sentence.

The next transducer 1909 resolves ambiguities in the choice of words for the target sentence arising from ambiguities in the process 1901 of generating adverbs and adjectives from their morphological units.

For example, the transducer 1909 converts the sentence

how quick\_quickly can you run?

which contains the ambiguity in resolving quick ADV ly into a single adverb, into the sentence

how quickly can you run?

To perform such conversions, the transducer 1909 uses a target-language model to assign a probability or score to each of the different possible sentences corresponding to an input sentence with a morphological ambiguity. The sentence with the highest probability or score is selected. In the above example, the sentence how quickly can you run? is selected because it has a higher target-language model probability or score than how quick can you run? In some embodiments of the transducer 1909 the target-language model is a trigram model similar to the target-structure language model 706. Such a transducer can be constructed by a person skilled in the art. The last transducer 1910 assigns a case to the words of a target sentence based on the casing rules for English. Principally this involves capitalizing the words at the beginning of sentences. Such a transducer can easily be constructed by a person skilled in the art.

### 5.3 TARGET TRANSDUCERS WITH CONSTRAINTS

In some embodiments, such as that depicted in FIG. 9, a target-structure transducer, such as that labelled 904, accepts a set of constraints that restricts its transformations of an intermediate target structure to target text.

Such constraints include, but are not limited to, requiring that a particular target word appear in the final target text. For example, because of ambiguities in a morphology, a target transducer may not be able to distinguish between electric and electrical. Thus the output of a target transducer might include the phrase electrical blanket in a situation where electric blanket is preferred. A constraint such as The word "electric" must appear in the target sentence would correct this shortcoming.

A target-transducer accepting such constraints in similar to target transducers as described in this section. Based on the descriptions already given, such transducers can be constructed by a person with a computer science background and skilled in the art.

## 6 TARGET LANGUAGE MODEL

The inventions described in this specification employ probabilistic models of the target language in a number of places. These include the target structure language model 705, and the class language model used by the decoder 404. As depicted in FIG. 20, the role of a language model is to compute an a priori probability or score of a target structure.

Language models are well known in the speech recognition art. They are described in the article "Self-Organized Language Modeling for Speech Recognition", by F. Jelinek, appearing in the book Readings in Speech Recognition edited by A. Waibel and K. F. Lee and published by Morgan Kaufmann Publishers, Inc., San Mateo, Calif. in 1990. They are also described in the article "A Tree-Based Statistical Model for Natural Language Speech Recognition", by L. Bahl, et al., appearing in the July 1989 Volume 37 of the IEEE Transactions on Acoustics, Speech and Signal Processing. These articles are included by reference herein. They are further described in the paper "Trainable Grammars for Speech Recognition", by J. Baker, appearing in the 1979 Proceedings of the Spring Conference of the Acoustical Society of America.

In some embodiments of the present inventions, the target structure consists of a sequence of morphs. In these embodiments, n-gram language models, as described in the aforementioned article by F. Jelinek, can be used. In other embodiments, the target structure comprises parse trees of the target language. In these embodiments, language models based on stochastic context-free grammars, as described in the aforementioned articles by F. Jelinek and the aforementioned paper by J. Baker, can be used.

In addition, decision tree language models, as described in the aforementioned paper by L. Bahl, et al. can be adapted by one skilled in the art to model a wide variety of target structures.

### 6.1 PERPLEXITY

The performance of a language model in a complete system depends on a delicate interplay between the language model and other components of the system. One language model may surpass another as part of a speech recognition system but perform less well in a translation system.

Since it is expensive to evaluate a language model in the context of a complete system, it is useful to have an intrinsic measure of the quality of a language model. One such measure is the probability that the model assigns to the large sample of target structures. One judges as better the language model which yields the greater probability. When the target structure is a sequence of words or morphs, this measure can be adjusted so that it takes account of the length of the structures. This leads to the notion of the perplexity of a language model with respect to a sample of text S:

$$\text{perplexity} = \Pr(S)^{-\frac{1}{|S|}} \quad (3)$$

where |S| is the number of morphs of S. Roughly speaking, the perplexity is the average number of morphs which the model cannot distinguish between, in predicting a morph of S. The language model with the smaller perplexity will be

the one which assigns the larger probability to S.

Because perplexity depends not only on the language model but also on the sample of text, it is important that the text be representative of that for which the language model is intended. Because perplexity is subject to sampling error, making fine distinctions between language models may require that the perplexity be measured with respect to a large sample.

### 6.2 n-GRAM LANGUAGE MODELS

n-gram language models will now be described. For these models, the target structure consists of a sequence of morphs.

Suppose  $m_1 m_2 m_3 \dots m_k$  be a sequence of  $k$  morphs  $m_i$ . For  $1 \leq i \leq j \leq k$ , let  $m_i^j$  denote the subsequence  $m_i m_{i+1} \dots m_j$ . For any sequence, the probability of a  $m_1^k$  is equal to the product of the conditional probabilities of each morph  $m_i$  given the previous morphs  $m_1^{i-1}$ :

$$Pr(m_1^k) = Pr(m_1)Pr(m_2|m_1)Pr(m_3|m_1m_2) \dots Pr(m_j|m_1^{j-1}) \dots Pr(m_k|m_1^{k-1}). \quad (20)$$

The sequence  $m_1^{i-1}$  is called the history of the morph  $m_i$  in the sequence.

For an n-gram model, the conditional probability of a morph in a sequence is assumed to depend on its history only through the previous  $n-1$  morphs:

$$Pr(m_i|m_1^{i-1}) = Pr(m_i|m_{i-n+1}^{i-1}) \quad (5)$$

For a vocabulary of size  $V$ , a 1-gram model is determined by  $V-1$  independent numbers, one probability  $Pr(m)$  for each morph  $m$  in the vocabulary, minus one for the constraint that all of the probabilities add up to 1. A 2-gram model is determined by  $V^2-1$  independent numbers,  $V(V-1)$  conditional probabilities of the form  $Pr(m_2|m_1)$  and  $V-1$  of the form  $Pr(m)$ . In general, an n-gram model is determined by  $V^n-1$  independent numbers,  $V^{n-1}(V-1)$  conditional probabilities of the form  $Pr(m_n|m_1^{n-1})$ , called the order-n conditional probabilities, plus  $V^{n-1}-1$  numbers which determine an (n-1) gram model.

The order-n conditional probabilities of an n-gram model form the transition matrix of an associated Markov model. The states of this Markov model are sequences of  $n-1$  morphs, and the probability of a transition from the state  $m_1 m_2 \dots m_{n-1}$  to the state  $m_2 m_3 \dots m_n$  is  $Pr(m_n|m_1 m_2 \dots m_{n-1})$ . An n-gram language model is called consistent if, for each string  $m_1^{n-1}$ , the probability that the model assigns to  $m_1^{n-1}$  is the steady state probability for the state  $m_1^{n-1}$  of the associated Markov model.

### 6.3 SIMPLE n-GRAM MODELS

The simplest form of an n-gram model is obtained by assuming that all the independent conditional probabilities are independent parameters. For such a model, values for the parameters can be determined from a large sample of training text by sequential maximum likelihood training. The order n-probabilities are given by

$$Pr(m_n|m_1^{n-1}) = \frac{f(m_1^{n-1} m_n)}{\sum_m f(m_1^{n-1} m)} \quad (6)$$

where  $f(m_1^j)$  is the number of times the string of morphs  $m_1^j$  appears in the training text. The remaining parameters are determined inductively by an analogous formula applied to the corresponding  $n-1$ -gram model. Sequential maximum likelihood training does not produce a consistent model, although for a large amount of training text, it produces a

model that is very nearly consistent.

Unfortunately, many of the parameters of a simple n-gram model will not be reliably estimated by this method. The problem is illustrated in Table 3, which shows the number of 1-2-, and 3-grams appearing with various frequencies in a sample of 365,893,263 words of English text from a variety of sources. The vocabulary consists of the 260,740 different words plus a special unknown word into which all other words are mapped. Of the  $6.799 \times 10^{10}$  2-grams that might have occurred in the data, only 14,494,217 actually did occur and of these, 8,045,024 occurred only once each. Similarly, of the  $1.733 \times 10^{16}$  3-grams that might have occurred, only 75,349,888 actually did occur and of these, 53,737,350 occurred only once each. These data and Turing's formula imply that 14.7 percent of the 3-grams and for 2.2 percent of the 2-grams in a new sample of English text will not appear in the original

TABLE 3

Number of n-grams with various frequencies in 365,893,263 words of running text.			
Count	1-grams	2-grams	3-grams
1	36,789	8,045,024	53,737,350
2	20,269	2,065,469	9,229,958
3	13,123	970,434	3,653,791
>3	135,335	3,413,290	8,728,789
>0	205,516	14,494,217	75,349,888
≥0	260,741	$6.799 \times 10^{10}$	$1.773 \times 10^{16}$

sample. Thus, although any 3-gram that does not appear in the original sample is rare, there are so many of them that their aggregate probability is substantial.

Thus, as  $n$  increases, the accuracy of a simple n-gram model increases, but the reliability of the estimates for its parameters decreases.

### 6.4 SMOOTHING

A solution to this difficulty is provided by interpolated estimation, which is described in detail in the paper "Interpolated estimation of Markov source parameters from sparse data", by F. Jelinek and R. Mercer and appearing in Proceedings of the Workshop on Pattern Recognition in Practice, published by North-Holland, Amsterdam, The Netherlands, in May 1980. Interpolated estimation combines several models into a smoothed model which uses the probabilities of the more accurate models where they are reliable and, where they are unreliable, falls back on the more reliable probabilities of less accurate models. If  $Pr^{(j)}(m_i|m_1^{i-1})$  is the  $j$ th language model, the smoothed model,  $Pr(m_i|m_1^{i-1})$ , is given by

$$Pr(m_i|m_1^{i-1}) = \sum_j \lambda_j(m_1^{i-1}) Pr^{(j)}(m_i|m_1^{i-1}). \quad (7)$$

The values of the  $\lambda_j(m_1^{i-1})$  are determined using the EM method, so as to maximize the probability of some additional sample of training text called held-out data. When interpolated estimation is used to combine simple 1-, 2-, and 3-gram models, the  $\lambda$ 's can be chosen to depend on  $m_1^{i-1}$  only through the count of  $m_{i-2} m_{i-1}$ . Where this count is high, the simple 3-gram model will be reliable, and, where this count is low, the simple 3-gram model will be unreliable.

The inventors constructed an interpolated 3-gram model in which the  $\lambda$ 's were divided into 1782 different sets according to the 2-gram counts, and determined from a held-out sample of 4,630,934 million words. The power of the model was tested using the 1,014,312 word Brown corpus. This well known corpus, which contains a wide

variety of English text, is described in the book Computational Analysis of Present-Day American English, by H. Kucera and W. Francis, published by Brown University Press, Providence, R.I., 1967. The Brown corpus was not included in either the training or held-out data used to construct the model. The perplexity of the interpolated model with respect to the Brown corpus was 244.

#### 6.5 n-GRAM CLASS MODELS

Clearly, some words are similar to other words in their meaning and syntactic function. For example, the probability distribution of words in the vicinity of Thursday is very much like that for words in the vicinity of Friday. Of course, they will not be identical: people rarely say Thank God it's Thursday! or worry about Thursday the 13<sup>th</sup>.

In class language models, morphs are grouped into classes, and morphs in the same class are viewed as similar. Suppose that  $\hat{c}$  is a map that partitions the vocabulary of  $V$  morphs into  $C$  classes by assigning each morph  $m$  to a class  $\hat{c}(m)$ . An  $n$ -gram class model based on  $\hat{c}$  is an  $n$ -gram language model for which

$$Pr(m_k | m_1^{k-1}) = Pr(m_k | c_k) Pr(c_k | c_1^{k-1}) \quad (8)$$

where  $c_i = \hat{c}(m_i)$ . An  $n$ -gram class model is determined by  $C^n - 1 + V - C$  independent numbers,  $V - C$  of the form  $Pr(m_i | c_i)$ , plus  $C^n - 1$  independent numbers which determine an  $n$ -gram language model for a vocabulary of size  $C$ . If  $C$  is much smaller than  $V$ , these are many fewer numbers than are required to specify a general  $n$ -gram language model.

In a simple  $n$ -gram class model, the  $C^n - 1 + V - C$  independent probabilities are treated as independent parameters. For such a model, values for the parameters can be determined by sequential maximum likelihood training. The order  $n$  probabilities are given by

$$Pr(c_n | c_1^{n-1}) = \frac{f(c_1^{n-1} c_n)}{\sum_c f(c_1^{n-1} c)} \quad (9)$$

where  $f(c_1^k)$  is the number of times that the sequence of classes  $c_1^k$  appears in the training text. (More precisely,  $f(c_1^k)$  is the number of distinct occurrences in the training text of a consecutive sequence of morphs  $m_1^k$  for which  $c_k = \hat{c}(m_k)$  for  $1 \leq k \leq i$ .)

### 7 CLASSES

The inventions described in this specification employ classes of morphs or words in a number of places. These include the class language model used by the decoder 702 and described in Section 14, and some embodiments of the target structure language model 705.

The inventors have devised a number of methods for automatically partitioning a vocabulary into classes based upon frequency or cooccurrence statistics or other information extracted from textual corpora or other sources. In this section, some of these methods will be explained. An application to construction of syntactic classes of words will be described. A person skilled in the art can easily adapt the methods to other situations. For example, the methods can be used to construct classes of morphs instead of classes of words. Similarly, they can be used to construct classes based upon cooccurrence statistics or statistics of word alignments in bilingual corpora.

#### 7.1 MAXIMUM MUTUAL INFORMATION CLUSTERING

A general scheme for clustering a vocabulary into classes is depicted schematically in FIG. 31. It takes as input a desired number of classes  $C$  3101, a vocabulary 3102 of size  $V$ , and a model 3103 for a probability distribution  $P(w_1, w_2)$  over bigrams from the vocabulary. It produces as output a partition 3104 of the vocabulary into  $C$  classes. In one application, the model 3103 can be a 2-gram language model as described in Section 6, in which case  $P(w_1, w_2)$  would be proportional to the number of times that the bigram  $w_1 w_2$  appears in a large corpus of training text.

Let the score  $\psi(\hat{C})$  of a partition  $\hat{C}$  be the average mutual information between the classes of  $\hat{C}$  with respect to the probability distribution  $P(w_1, w_2)$ :

$$\psi(\hat{C}) = \sum_{c_1, c_2} P(c_1, c_2) \log \frac{P(c_1, c_2)}{P(c_1)P(c_2)} \quad (10)$$

In this sum,  $c_1$  and  $c_2$  each run over the classes of the partition  $\hat{C}$ , and

$$P(c_1, c_2) = \sum_{w_1 \in c_1, w_2 \in c_2} P(w_1, w_2) \quad (11)$$

$$P(c_1) = \sum_{c_2} P(c_1, c_2) \quad (12)$$

$$P(c_1) = \sum_{c_2} P(c_1, c_2) \quad (13)$$

The scheme of FIG. 31 chooses a partition  $\hat{C}$  for which the score average mutual information  $\psi(\hat{C})$  is large.

#### 7.2 A CLUSTERING METHOD

One method 3204 for carrying out this scheme is depicted in FIG. 32. The method proceeds iteratively. It begins (step 3203 with a partition of size  $V$  in which each word is assigned to a distinct class. At each stage of the iteration (Steps 3201 and 3202), the current partition is replaced by a new partition which is obtained by merging a pair of classes into a single class. The pair of classes to be merged is chosen so that the score of the new partition is as large as possible. The method terminates after  $V - C$  iterations, at which point the current partition contains  $C$  classes.

In order that it be practical, the method 3204 must be implemented carefully. At the  $i^{\text{th}}$  iteration, a pair of classes to be merged must be selected from amongst approximately  $(V-i)^2/2$  pairs. The score of the partition obtained by merging any particular pair is the sum of  $(V-i)^2$  terms, each of which involves a logarithm. Since altogether there are  $V - C$  merges, this straight-forward approach to the computation is of order  $V^5$ . This is infeasible, except for very small values of  $V$ . A more frugal organization of the computation must take advantage of the redundancy in this straight-forward calculation.

An implementation will now be described in which the method 3204 executes in time of order  $V^3$ . In this implementation, the change in score due to a merge is computed in constant time, independent of  $V$ .

Let  $\hat{C}_k$  denote the partition after  $V - k$  merges. Let  $\hat{C}_k(1), \hat{C}_k(2), \dots, \hat{C}_k(k)$  denote the  $k$  classes of  $\hat{C}_k$ . Let  $p_k(l, m) = P(\hat{C}_k(l), \hat{C}_k(m))$  and let

$$p_k(l) = \sum_m p_k(l, m) \quad (14)$$



-continued

$$pr_k(m) = \sum_l p_k(l, m) \quad (15)$$

$$q_k(l, m) = p_k(l, m) \log \frac{p_k(l, m)}{p_k(l)pr_k(m)} \quad (16)$$

Let  $I_k = \psi(\hat{C}_k)$  be the score of  $\hat{C}_k$ , so that

$$I_k = \sum_{l, m} q_k(l, m) \quad (17)$$

Let  $I_k(i, j)$  be the score of the partition obtained from  $\hat{C}_k$  by merging classes  $\hat{C}_k(i)$  and  $\hat{C}_k(j)$ , and let  $L_k(i, j) = I_k - I_k(i, j)$  be the change in score as a result of this merge. Then

$$L_k(i, j) = s_k(i) + s_k(j) - q_k(i, j) - q_k(j, i) - q_k(i \cup j, i \cup j) - \quad (18)$$

$$\sum_{l \neq i, j} q_k(l, i \cup j) - \sum_{m \neq i, j} q_k(i \cup j, m), \quad (18)$$

where

$$s_k(i) = \sum_l q_k(l, i) + \sum_m q_k(i, m). \quad (19)$$

In these and subsequent formulae,  $i \cup j$  denotes the result of the merge, so that, for example

$$p_k(i \cup j, m) = p_k(l, m) + p_k(j, m) \quad (20)$$

$$q_k(i \cup j, m) = p_k(i \cup j, m) \log \frac{p_k(i \cup j, m)}{p_k(i \cup j)pr_k(m)} \quad (21)$$

The key to the implementation is to store and inductively update the quantities

$$\begin{array}{ccccccc} p_k(l, m) & p_k(l) & pr_k(m) & q_k(l, m) & & & \\ I_k & s_k(i) & L_k(i, j) & & & & \end{array} \quad (22)$$

Note that if  $I_k$ ,  $s_k(i)$ , and  $s_k(j)$  are known, then the majority of the time involved in computing  $I_k(i, j)$  is devoted to computing the sums on the second line of equation 18. Each of these sums has approximately  $V-k$  terms and so this reduces the problem of evaluating  $I_k(i, j)$  from one of order  $V^2$  to one of order  $V$ .

Suppose that the quantities shown in Equation 22 are known at the beginning of an iteration. Then the new partition  $\hat{C}_{k+1}$  is obtained by merging the pair of classes  $\hat{C}_k(i)$  and  $\hat{C}_k(j)$ ,  $i < j$ , for which  $L_k(i, j)$  is smallest. The  $k-1$  classes of the new partition are  $\hat{C}_{k-1}(1)$ ,  $\hat{C}_{k-1}(2)$ , ...,  $\hat{C}_{k-1}(k-1)$  with

$$C_{k-1}(l) = C_k(l) \text{ if } l \neq i, j$$

$$C_{k-1}(i) = C_k(i) \cup C_k(j)$$

$$C_{k-1}(j) = C_k(j) \text{ if } j < k$$

Obviously,  $I_{k+1} = I_k(i, j)$ . The values of  $p_{k+1}$ ,  $p_{k+1}$ ,  $pr_{k+1}$ , and  $q_{k+1}$  can be obtained easily from  $p_k$ ,  $p_k$ ,  $pr_k$ , and  $q_k$ . If  $l$  and  $m$  denote indices neither of which is equal to either  $i$  or  $j$ , then

$$s_{k+1}(l) = s_k(l) - q_k(l, i) - q_k(l, j) - q_k(i, l) - q_k(j, l) + q_{k+1}(l, i) + q_{k+1}(l, j)$$

$$s_{k+1}(j) = s_k(j) - q_k(j, i) - q_k(i, j) - q_k(j, i) - q_k(i, j) + q_{k+1}(j, i) + q_{k+1}(j, j)$$

$$L_{k+1}(l, m) = L_k(l, m) - q_k(l, m, i) - q_k(l, m, j) - q_k(l, m, i) - q_k(l, m, j) + q_{k+1}(l, m, i) + q_{k+1}(l, m, j)$$

$$L_{k+1}(i, j) = L_k(i, j) - q_k(i, j, i) - q_k(i, j, j) - q_k(i, j, i) - q_k(i, j, j) + q_{k+1}(i, j, i) + q_{k+1}(i, j, j)$$

$$i) + q_{k+1}(i, j)$$

$$L_{k+1}(j, i) = L_{k+1}(i, j) \quad (23)$$

Finally,  $s_{k+1}(i)$  and  $L_{k+1}(l, i)$  are determined from equations 18 and 19.

This update process requires order  $V^2$  computations. Thus, by this implementation, each iteration of the method requires order  $V^2$  time, and the complete method requires order  $V^3$  time.

The implementation can be improved further by keeping track of those pairs  $l, m$  for which  $p_k(l, m)$  is different from zero. For example, suppose that  $P$  is given by a simple bigram model trained on the data described in Table 3 of Section 6. In this case, of the  $6.799 \times 10^{10}$  possible word 2-grams  $w_1, w_2$ , only 14,494,217 have non-zero probability. Thus, in this case, the sums required in equation 18 have, on average, only about 56 non-zero terms instead of 260,741 as might be expected from the size of the vocabulary.

### 7.3 AN ALTERNATE CLUSTERING METHOD

For very large vocabularies, the method 3204 may be too computationally costly. The following alternate method can be used. First, the words of the vocabulary are arranged in order of frequency with the most frequent words first. Each of the first  $C$  words is assigned to its own distinct class. The method then proceeds iteratively through  $V-C$  steps. At the  $k^{\text{th}}$  step the  $(C+k)^{\text{th}}$  most probable word is assigned to a new class. Then, two of the resulting  $C+1$  classes are merged into a single class. The pair of classes that is merged is the one for which the loss is average mutual information is least. After  $V-C$  steps, each of the words in the vocabulary will have been assigned to one of  $C$  classes.

### 7.4 IMPROVING CLASSES

The classes constructed by the clustering method 3204 or the alternate clustering method described above can often be improved. One method of doing this is to repeatedly cycle through the vocabulary, moving each word to the class for which the resulting partition has the highest average mutual information score. Eventually, no word will move and the method finishes. It may be possible to further improve the classes by simultaneously moving two or more words, but for large vocabularies, such a search is too costly to be feasible.

### 7.5 EXAMPLES

The methods described above were used to divide the 260,741-word vocabulary of Table 3, Section 6, into 1000 classes. Table 4 shows some of the classes that are particularly interesting, and Table 5 shows classes that were selected at random. Each of the lines in the tables contains members of a different class. The average class has 260 words. The table shows only those words that occur at least ten times, and only the ten most frequent words of any class. (The other two months would appear with the class of months if this limit had been extended to twelve). The degree to which the classes capture both syntactic and semantic aspects of English is quite surprising given that they were constructed from nothing more than counts of bigrams. The class {that tha theat} is interesting because although tha and theat are English words, the method has discovered that in the training data each of them is most often a mistyped that.

### 7.6 A METHOD FOR CONSTRUCTING SIMILARITY TREES

The clustering method 3204 can also be used to construct a similarity tree over the vocabulary. Suppose the merging steps 3201 and 3202 of method 3204 are iterated  $V-1$  times, resulting in a single class consisting of the entire vocabulary. The order in which the classes are merged determines a binary tree, the root of which corresponds to this single class



and the leaves of which correspond to the words in the vocabulary. Intermediate nodes of the tree correspond to groupings of words intermediate between single words and the entire vocabulary. Words that are statistically similar with respect to the model  $P(w_1, w_2)$  will be close together in the tree.

FIG. 30 shows some of the substructures in a tree constructed in this manner using a simple 2-gram model for the 1000 most frequent words in a collection of office correspondence.

## 8 OVERVIEW OF TRANSLATION MODELS AND PARAMETER ESTIMATION

This section is an introduction to translation models and methods for estimating the parameters of such models. A more detailed discussion of these topics will be given in Section 9.

TABLE 4

Classes from a 260,741-word vocabulary

Friday Monday Thursday Wednesday Tuesday Saturday Sunday  
weekends Sundays Saturdays  
June March July April January December October November  
September August  
people guys folks fellows CEOs chaps doubters commies un-  
fortunates blokes  
down backwards ashore sideways southward northward over-  
board aloft downwards adrift  
water gas coal liquid acid sand carbon steam shale iron  
great big vast sudden mere sheer gigantic lifelong scant colossal  
man woman boy girl lawyer doctor guy farmer teacher citizen  
American Indian European Japanese German African Catholic  
Israeli Italian Arab  
pressure temperature permeability density porosity stress velocity  
viscosity gravity tension  
mother wife father son husband brother daughter sister boss uncle  
machine device controller processor CPU printer spindle sub-  
system compiler plotter  
John George James Bob Robert Paul William Jim David Mike  
anyone someone anybody somebody  
feet miles pounds degrees inches barrels tons acres meters bytes  
director chief professor commissioner commander treasurer  
founder superintendent dean custodian  
liberal conservative parliamentary royal progressive Tory pro-  
visional separatist federalist PQ  
hadn't had would've could've should've must've might've  
asking telling wondering instructing informing kidding reminding  
bothering thanking deposing  
that the threat  
head body hands eyes voice arm seat eye hair mouth

TABLE 5

Randomly selected word classes

little prima moment's trifle tad Little minute's tinker's hornet's  
teammate's  
ask remind instruct urge interrupt invite congratulate commend  
warn applaud  
object apologize apologise avow wish  
cost expense risk profitability deferral earmarks capstone  
cardinality mintage reseller  
B dept. AA Whitey CL pi Namerow PA Mgr. LaRose  
# Rel rel. #S Shree  
S Gens nai Matsuzawa ow Kageyama Nishida Sumit Zollner  
Mallik  
research training education science advertising arts medicine  
machinery Art AIDS  
rise focus depend rely concentrate dwell capitalize embark intrude  
typewriting  
Minister mover Sydneys Minister Minister  
running moving playing setting holding carrying passing cutting  
driving fighting

TABLE 5-continued

Randomly selected word classes

court judge jury slam Edelstein magistrate marshal Abella Scalia  
larceny  
annual regular monthly daily weekly quarterly periodic Good  
yearly convertible  
aware unaware unsure cognizant apprised mindful partakers  
force ethic stoppage force's conditioner stoppages conditioners  
waybill forwarder Atonabee  
systems magnetics loggers products' coupler Econ databanks  
Centre inscriber correctors  
industry producers makers fishery Arabia growers addiction  
medalist inhalation addict  
brought moved opened picked caught tied gathered cleared hung  
lifted

### 8.1 TRANSLATION MODELS

As illustrated in FIG. 21, a target structure to source structure translation model  $P_\theta$  706 with parameters  $\theta$  is a method for calculating a conditional probability, or likelihood,  $P_\theta(f|e)$ , for any source structure  $f$  given any target structure  $e$ . Examples of such structures include, but are not limited to, sequences of words, sequences of linguistic morphs, parse trees, and case frames. The probabilities satisfy:

$$P_\theta(f|e) \geq 0, P_\theta(\text{failure}|e) \geq 0, \quad (24)$$

$$P_\theta(\text{failure}|e) + \sum_f P_\theta(f|e) = 1,$$

where the sum ranges over all structures  $f$ , and failure is a special symbol.  $P_\theta(f|e)$  can be interpreted as the probability that a translator will produce  $f$  when given  $e$ , and  $P_\theta(\text{failure}|e)$  can be interpreted as the probability that he will produce no translation when given  $e$ . A model is called deficient if  $P_\theta(\text{failure}|e)$  is greater than zero for some  $e$ .

#### 8.1.1 Training

Before a translation model can be used, values must be determined for the parameters  $\theta$ . This process is called parameter estimation or training.

One training methodology is maximum likelihood training, in which the parameter values are chosen so as to maximize the probability that the model assigns to a training sample consisting of a large number  $S$  of translations  $(f^{(s)}, e^{(s)})$ ,  $s=1, 2, \dots, S$ . This is equivalent to maximizing the log likelihood objective function

$$\psi(P_\theta) = S^{-1} \sum_{s=1}^S \log P_\theta(f^{(s)}|e^{(s)}) = \sum_{f,e} C(f,e) \log P_\theta(f|e). \quad (25)$$

Here  $C(f,e)$  is the empirical distribution of the sample, so that  $C(f,e)$  is  $1/S$  times the number of times (usually 0 or 1) that the pair  $(f,e)$  occurs in the sample.

#### 8.1.2 Hidden Alignment Models

In some embodiments, translation models are based on the notion of an alignment between a target structure and a source structure. An alignment is a set of connections between the entries of the two structures.

FIG. 22 shows an alignment between the target structure John V past 3s to kiss Mary and the source structure Jean V past 3s embrasser Marie. (These structures are transduced versions of the sentences John kissed Mary and Jean a embrassé Marie respectively.) In this alignment, the entry John is connected with the entry Jean, the entry V past 3s is connected with the entry V past 3s, the entry to kiss is connected with the entry embrasser, and the entry Mary is

connected with the entry Marie. FIG. 23 shows an alignment between the target structure John V<sub>past</sub> 3s to kiss the girl and the source structure Le jeune fille V<sub>passive</sub> 3s embrasser par Jean. (These structures are transduced versions of the sentences John kissed the girl and Le jeune fille a été embrassée par Jean respectively.)

In some embodiments, illustrated in FIG. 24, a translation model 706 computes the probability of a source structure given a target structure as the sum of the probabilities of all alignments between these structures:

$$P_{\theta}(fle) = \sum_{a} P_{\theta}(f, a | e). \quad (26)$$

In some embodiments, a translation model 706 can compute the probability of a source structure given a target structure as the maximum of the probabilities of all alignments between these structures:

$$P_{\theta}(fle) = \max_a P_{\theta}(f, a | e). \quad (27)$$

As depicted in FIG. 25, the probability  $P_{\theta}(fle)$  of a single alignment is computed by a detailed translation model 2101. The detailed translation model 2101 employs a table 2501 of values for the parameters  $\theta$ .

## 8.2 AN EXAMPLE OF A DETAILED TRANSLATION MODEL

A simple embodiment of a detailed translation model is depicted in FIG. 26. This embodiment will now be described. Other, more sophisticated, embodiments will be described in in Section 9.

For the simple embodiment of a detailed translation model depicted in FIG. 26, the source and target structures are sequences of linguistic morphs. This embodiment comprises three components:

1. a fertility sub-model 2601;
2. a lexical sub-model 2602;
3. a distortion sub-model 2603.

The probability of an alignment and source structure given a target structure is obtained by combining the probabilities computed by each of these sub-models. Corresponding to these sub-models, the table of parameter values 2501 comprises:

- 1a. fertility probabilities  $n(\phi | e)$  where  $\phi$  is any non-negative integer and  $e$  is any target morph;
- b. null fertility probabilities  $n_0(\phi | m')$ , where  $\phi$  is any non-negative integer and  $m'$  is any positive integer;
- 2a. lexical probabilities  $t(f | e)$ , where  $f$  is any source morph, and  $e$  is any target morph;
- b. lexical probabilities  $t(f | \text{*null*})$ , where  $f$  is any source morph, and  $\text{*null*}$  is a special symbol;
3. distortion probabilities  $a(j | i, m)$ , where  $m$  is any positive integer,  $i$  is any positive integer, and  $j$  is any positive integer between 1 and  $m$ .

This embodiment of the detailed translation model 2101 computes the probability  $P_{\theta}(f, a | e)$  of an alignment  $a$  and a source structure  $f$  given a target structure  $e$  as follows. If any source entry is connected to more than one target entry, then the probability is zero. Otherwise, the probability is computed by the formula

$$P_{\theta}(f, a | e) = \text{combinatorial\_factor} \times \text{fertility\_prob} \times \text{lexical\_prob} \times \text{distortion\_prob} \quad (28)$$

where

$$\text{combinatorial\_factor} = \frac{1}{\pi} \prod_{i=0}^l (\phi_i!) \quad (29)$$

$$\text{fertility\_prob} = n_0 \left( \phi_0! \prod_{i=1}^l \phi_i! \right) \prod_{i=1}^l n(\phi_i | e_i) \quad (30)$$

$$\text{lexical\_prob} = \prod_{j=0}^m t(f_j | e_{a_j}) \quad (31)$$

$$\text{distortion\_prob} = \frac{1}{\phi_0!} \prod_{j: a_j \neq 0} a(j | a_j, m) \quad (32)$$

Here

- $l$  is the number of morphs in the target structure;
- $e_i$  for  $i=1, 2, \dots, l$  is the  $i^{\text{th}}$  morph of the target structure;
- $e_0$  is the special symbol  $\text{*null*}$ ;
- $m$  is the number of morphs in the source structure;
- $f_j$  for  $j=1, 2, \dots, m$  is the  $j^{\text{th}}$  morph of the source structure;
- $\phi_0$  is the number of morphs of the target structure that are not connected with any morphs of the source structure;
- $\phi_i$  for  $i=1, 2, \dots, l$  is the number of morphs of the target structure that are connected with the  $i^{\text{th}}$  morph of the source structure;
- $a_j$  for  $j=1, 2, \dots, m$  is 0 if the  $j^{\text{th}}$  morph of the source structure is not connected to any morph of the target structure;
- $a_j$  for  $j=1, 2, \dots, m$  is  $i$  if the  $j^{\text{th}}$  morph of the source structure is connected to the  $i^{\text{th}}$  morph of the target structure.

These formulac are best illustrated by an example. Consider the source structure, target structure, and alignment depicted in FIG. 23. There are 7 source morphs, so  $m=7$ . The first source morph is Le, so  $f_1=\text{Le}$ ; the second source morph is jeune, so  $f_2=\text{jeune}$ ; etc. There are 5 target morphs so  $l=5$ . The first target morph is John, so  $e_1=\text{John}$ ; the second target morph is V<sub>past</sub> 3s, so  $e_2=\text{V past 3s}$ ; etc. The first source morph (Le) is connected to the fourth target morph (the), so  $a_1=4$ . The second source morph (jeune) is connected to the fifth target morph (girl) so  $a_2=5$ . The complete alignment is

$$a_1 = 4 \quad a_2 = 5 \quad a_3 = 5 \quad a_4 = 2 \quad a_5 = 3 \quad a_6 = 0 \quad a_7 = 1$$

All target morphs are connected to at least one source morph, so  $\phi_0=0$ . The first target morph (John) is connected to exactly one source morph (Jean), so  $\phi_1=1$ . The second target morph (V<sub>past</sub> 3s) is connected to exactly one source morph (V<sub>passive</sub> 3s), so  $\phi_2=1$ . The complete set of fertilities are

$$\phi_0 = 0 \quad \phi_1 = 1 \quad \phi_2 = 1 \quad \phi_3 = 1 \quad \phi_4 = 1 \quad \phi_5 = 2$$

The components of the probability  $P_{\theta}(f, a | e)$  are

$$\text{combinatorial\_factor} = 0!1!1!1!1!2! \quad \text{note: } 0! = 1$$

-continued

fertility\_prob = no(k=15)  
 $n(1|John)n(1|V\_past\_3s)$   
 $n(1|to\_kiss)n(1|he)n(1|girl)$

lexical\_prob =  $n(Le\ he)n(juncelgirl)n(fillc|girl)$   
 $n(V\_passive\_3s|V\_past\_3s)n(embrasser|to\_kiss)$   
 $n(part*null*)n(Jean|John)$

distortion\_prob =  $\frac{1}{0!} a(1|4,7)a(2|5,7)a(3|5,7)$   
 $a(4|2,7)a(5|3,7)a(7|1,7)$

Many other embodiments of the detailed translation 2101 are possible. Five different embodiments will be described in Section 9 and Section 10.

### 8.3 ITERATIVE IMPROVEMENT

It is not possible to create of sophisticated model or find good parameter values at a stroke. One aspect of the present invention is a method of model creation and parameter estimation that involves gradual iterative improvement. For a given model, current parameter values are used to find better ones. In this way, starting from initial parameter values, locally optimal values are obtained. Then, good parameter values for one model are used to find initial parameter values for another model. By alternating between these two steps, the method proceeds through a sequence of gradually more sophisticated models.

As illustrated in FIG. 27, the method comprises the following steps:

2701. Start with a list of models Model 0, Model 1, Model 2 etc.
2703. Choose parameter values for Model 0.
2702. Let  $i=0$ .
2707. Repeat steps 2704, 2705 and 2706 until all models are processed.
2704. Increment  $i$  by 1.
2705. Use Models 0 to  $i-1$  together with their parameter values to find initial parameter values for Model  $i$ .
2706. Use the initial parameter values for Model  $i$ , to find locally optimal parameter values for Model  $i$ .

#### 8.3.1 Relative Objective Function

Two hidden alignment models  $\tilde{P}_\theta$  and  $P_\theta$  of the form depicted in FIG. 24 can be compared using the relative objective function<sup>1</sup>

$$R(P_\theta, \tilde{P}_\theta) = \sum_{f,e} C(f,e) \sum_a P_\theta(a|f,e) \log \frac{P_\theta(f,ale)}{P_\theta(f,ale)} \quad (34)$$

where  $\tilde{P}_\theta(a|f,e) = \tilde{P}_\theta(a, f|e) / \tilde{P}_\theta(f|e)$ . Note that  $R(\tilde{P}_\theta, \tilde{P}_\theta) = 0$ .  $R$  is related to  $\psi$  by Jensen's inequality

$$\psi(P_\theta) \geq \psi(\tilde{P}_\theta) + R(P_\theta, \tilde{P}_\theta), \quad (35)$$

which follows because the logarithm is concave. In fact, for any  $e$  and  $f$ ,

$$\sum_a P_\theta(a|f,e) \log \frac{P_\theta(f,ale)}{P_\theta(f,ale)} \geq \log \sum_a P_\theta(a|f,e) \frac{P_\theta(f,ale)}{P_\theta(f,ale)} \quad (36)$$

$$= \log \frac{P_\theta(f|e)}{P_\theta(f|e)} = \log P_\theta(f|e) - \log P_\theta(f|e). \quad (37)$$

$$D(p,q) = \sum_x p(x) \log \frac{p(x)}{q(x)} \quad (33)$$

between probability distributions  $p$  and  $q$ . However, whereas the relative entropy is never negative,  $R$  can take any value. The inequality (35) for  $R$  is the analog of the inequality  $D \geq 0$  for  $D$ .

#### 8.3.2 Improving Parameter Values

From Jensen's inequality (35), it follows that  $\psi(P_\theta)$  is greater than  $\psi(\tilde{P}_\theta)$  if  $R(\tilde{P}_\theta, P_\theta)$  is positive. With  $\tilde{P} = P$ , this suggests the following iterative procedure, depicted in FIG. 28 and known as the EM Method, for finding locally optimal parameter values  $\theta$  for a model  $P$ :

2801. Choose initial parameter values  $\theta$ .
2803. Repeat steps 2802–2805 until convergence.
2802. With  $\theta$  fixed, find the values  $\theta$  which maximize  $R(P_\theta, P_\theta)$ .
2805. Replace  $\theta$  by  $\theta$ .
2804. After convergence,  $\theta$  will be locally optimal parameter values for model  $P$ .

Note that for any  $\theta$ ,  $R(P_\theta, P_\theta)$  is non-negative at its maximum in  $\theta$ , since it is zero for  $\theta = \theta$ . Thus  $\psi(P_\theta)$  will not decrease from one iteration to the next.

#### 8.3.3 Going from One Model to Another

Jensen's inequality also suggests a method for using parameter values  $\theta$  for one model  $\tilde{P}$  to find initial parameter values  $\theta$  for another model  $P$ :

2804. Start with parameter values  $\theta$  for  $\tilde{P}$
  2901. With  $\tilde{P}$  and  $\theta$  fixed, find the values  $\theta$  which maximize  $R(\tilde{P}_\theta, P_\theta)$ .
  2902.  $\theta$  will be good initial values for model  $P$ .
- In contrast to the case where  $\tilde{P} = P$ , there may not be any  $\theta$  for which  $R(\tilde{P}_\theta, P_\theta)$  is non-negative. Thus, it could be that, even for the best  $\theta$ ,  $\psi(P_\theta) < \psi(P_\theta)$ .

#### 8.3.4 Parameter Reestimation Formulae

To apply these procedures, it is necessary to solve the maximization problems of Steps 2802 and 2901. For the models described below, this can be done explicitly. To see the basic form of the solution, suppose  $P_\theta$  is a simple model given by

$$P_\theta(f,ale) = \pi \theta(\omega)^{c(\omega;a,f,e)}, \quad (38)$$

where the  $\theta(\omega), \omega \in \Omega$ , are real-valued parameters satisfying the constraints

$$\theta(\omega) \geq 0, \quad \sum_{\omega \in \Omega} \theta(\omega) = 1 \quad (39)$$

and for each  $\omega$  and  $(a,f,e)$ ,  $c(\omega;a,f,e)$  is a non-negative integer.<sup>2</sup> It is natural to interpret  $\theta(\omega)$  as the probability of the event  $\omega$  and  $c(\omega;a,f,e)$  as the number of times that this event occurs in  $(a,f,e)$ . Note that

<sup>2</sup>More generally, we can allow  $c(\omega;a,f,e)$  to be a non-negative real number.

$$c(\omega, f, e) = \theta(\omega) \frac{\partial}{\partial \theta(\omega)} \log P_{\theta}(f, a|e). \quad (40)$$

The values for  $\theta$  that maximize the relative objective function  $R(\bar{P}_{\theta}, P_{\theta})$  subject to the constraints (39) are determined by the Kuhn-Tucker conditions

$$\frac{\partial}{\partial \theta(\omega)} R(P_{\theta}, P_{\theta}) - \lambda = 0, \quad \omega \in \Omega, \quad (41)$$

where  $\lambda$  is a Lagrange multiplier, the value of which is determined by the equality constraint in Equation (39). These conditions are both necessary and sufficient for a maximum since  $R(\bar{P}_{\theta}, P_{\theta})$  is a concave function of the  $\theta(\omega)$ . Multiplying Equation (41) by  $\theta(\omega)$  and using Equation (40) and Definition (34) of  $R$ , yields the parameter reestimation formulae

$$\theta(\omega) = \lambda^{-1} c_{\theta}(\omega), \quad \lambda = \sum_{\omega \in \Omega} c_{\theta}(\omega), \quad (42)$$

$$c_{\theta}(\omega) = \sum_{f, e} C(f, e) c_{\theta}(\omega, f, e), \quad (43)$$

$$c_{\theta}(\omega, f, e) = \sum_a P_{\theta}(a|f, e) c(\omega, a, f, e). \quad (44)$$

$c_{\theta}(\omega, f, e)$  can be interpreted as the expected number of times, as computed by the model  $\bar{P}_{\theta}$ , that the event  $\omega$  occurs in the translation of  $e$  to  $f$ . Thus  $\theta(\omega)$  is the (normalized) expected number of times, as computed by model  $\bar{P}_{\theta}$ , that  $\omega$  occurs in the translations of the training sample.

These formulae can easily be generalized to models of the form (38) for which the single equality constraint in Equation (39) is replaced by multiple constraints

$$\sum_{\omega \in \Omega_{\mu}} \theta(\omega) = 1, \quad \mu = 1, 2, \dots, \quad (45)$$

where the subsets  $\Omega_{\mu}$ ,  $\mu = 1, 2, \dots$  form a partition of  $\Omega$ . Only Equation (42) needs to be modified to include a different  $\lambda_{\mu}$  for each  $\mu$ ; if  $\omega \in \Omega_{\mu}$ , then

$$\theta(\omega) = \lambda_{\mu}^{-1} c_{\theta}(\omega), \quad \lambda_{\mu} = \sum_{\omega \in \Omega_{\mu}} c_{\theta}(\omega). \quad (46)$$

### 8.3.5 Approximate and Viterbi Parameter Estimation

The computation required for the evaluation of the counts  $\bar{c}$  in Equation 44 increases with the complexity of the model whose parameters are being determined.

For simple models, such as Model 1 and Model 2 described below, it is possible to calculate these counts exactly by including the contribution of each possible alignments. For more sophisticated models, such as Model 3, Model 4, and Model 5 described below, the sum over alignments in Equation 44 is too costly to compute exactly. Rather, only the contributions from a subset of alignments can be practically included. If the alignments in this subset account for most of the probability of a translation, then this truncated sum can still be a good approximation.

This suggests the following modification to Step 2802 of the iterative procedure of FIG. 28:

- In calculating the counts using the update formulae 44, approximate the sum by including only the contributions from some subset of alignments of high probability.

If only the contribution from the single most probable alignment is included, the resulting procedure is called Viterbi Parameter Estimation. The most probable alignment between a target structure and a source structure is called the Viterbi alignment. The convergence of Viterbi Estimation is easily demonstrated. At each iteration, the parameter values are re-estimated so as to make the current set of Viterbi alignments as probable as possible; when these parameters are used to compute a new set of Viterbi alignments, either the old set is recovered or a set which is yet more probable is found. Since the probability can never be greater than one, this process surely converge. In fact, it converge in a finite, though very large, number of steps because there are only a finite number of possible alignments for any particular translation.

In practice, it is often not possible to find Viterbi alignments exactly. Instead, an approximate Viterbi alignment can be found using a practical search procedure.

### 8.4 Five Models

Five detailed translation models of increasing sophistication will be described in Section 9 and Section 10. These models will be referred to as Model 1, Model 2, Model 3, Model 4, and Model 5.

Model 1 is very simple but it is useful because its likelihood function is concave and consequently has a global maximum which can be found by the EM procedure. Model 2 is a slight generalization of Model 1. For both Model 1 and Model 2, the sum over alignments for the objective function and the relative objective function can be computed very efficiently. This significantly reduces the computational complexity of training. Model 3 is more complicated and is designed to more accurately model the relation between a morph of  $e$  and the set of morphs in  $f$  to which it is connected. Model 4 is a more sophisticated step in this direction. Both model 3 and Model 4 are deficient. Model 5 is a generalization of Model 4 in this deficiency is removed at the expense of more increased complexity. For Models 3, 4 and 5 the exact sum over alignments can not be computed efficiently. Instead, this sum can be approximated by restricting it to alignments of high probability.

Model 5 is a preferred embodiment of a detailed translation model. It is a powerful but unwieldy ally in the battle to describe translations. It must be led to the battlefield by its weaker but more agile brethren Models 2, 3, and 4. In fact, this is the *raison d'être* of these models.

## 9 DETAILED DESCRIPTION OF TRANSLATION MODELS AND PARAMETER ESTIMATION

In this section embodiments of the statistical translation model that assigns a conditional probability to the event that a sequence of lexical units in the source language is a translation of a sequence of lexical units in the target language will be described. Methods for estimating the parameters of these embodiments will be explained.

For concreteness the discussion will be phrased in terms of a source language of French and a target language of English. The sequences of lexical units will be restricted to sequences of words.

It should be understood that the translation models, and the methods for their training generalize easily to other source and target languages, and to sequences of lexical units comprising lexical morphemes, syntactic markers, sense labels, case frame labels, etc.

## 9.1 Notation

Random variables will be denoted by upper case letters, and the values of such variables will be denoted by the corresponding lower case letters. For random variables  $X$  and  $Y$ , the probability  $\Pr(Y=y|X=x)$  will be denoted by the shorthand  $P(y|x)$ . Strings or vectors will be denoted by bold face letters, and their entries will be denoted by the corresponding non-bold letters.

In particular,  $e$  will denote an English word,  $e$  will denote a string of English words, and  $E$  will denote a random variable that takes as values strings of English words. The length of an English string will be denoted  $l$ , and the corresponding random variable will be denoted  $L$ . Similarly,  $f$  will denote a French word,  $f$  will denote a string of English words, and  $F$  will denote a random variable that takes as values strings of French words. The length of a French string will be denoted  $m$ , and the corresponding random variables will be denoted  $M$ .

A table notation appears in Section 10.

## 9.2 Translations

A pair of strings that are translations of one another will be called a translation. A translation will be depicted by enclosing the strings that make it up in parentheses and separating them by a vertical bar. Thus,  $(\text{Qu'aurons-nous pu faire?}|\text{What could we have done?})$  depicts the translation of  $\text{Qu'aurons-nous pu faire?}$  as  $\text{What could we have done?}$ . When the strings are sentences, the final stop will be omitted unless it is a question mark or an exclamation point.

## 9.3 Alignments

Some embodiments of the translation model involve the idea of an alignment between a pair of strings. An alignment consists of a set of connections between words in the French string and words in the English string. An alignment is a random variable,  $A$ ; a generic values of this variable will be denoted by  $a$ . Alignments are shown graphically, as in FIG. 33, by writing the English string above the French string and drawing lines from some of the words in the English string to some of the words in the French string. These lines are called connections. The alignment in FIG. 33 has seven connections, (the, Le), (program, programme), and so on. In the description that follows, such alignment will be denoted as (Le programme a été mis en application|And the(1) program(2) has(3)been(4) implemented(5,6,7)). The list of numbers following an English word shows the positions in the French string of the words with which it is aligned. Because And is not aligned with any French words here, there is no list of numbers after it. Every alignment is considered to be correct with some probability. Thus (Le programme a été mis en application|And(1,2,3,4,5,6,7) the program has been implemented) is perfectly acceptable. Of course, this is much less probable than the alignment shown in FIG. 33.

As another example, the alignment (Le reste appartenait aux autochtones|The(1) balance(2) was(3) the(3) territory(3) of(4) the(4) aboriginal(5) people(5)) is shown in FIG. 34. Sometimes, two or more English words may be connected to two or more French words as shown in FIG. 35. This will be denoted as (Les pauvres sont démunis|The(1) poor(2) don't(3,4) have(3,4) any(3,4) money(3,4). Here, the four English words don't have any money conspire to generate the two French words sont démunis.

The set of English words connected to a French word will be called the notion that generates it. An alignment resolves the English string into a set of possibly overlapping notions that is called a notional scheme. The previous example contains the three notions. The, poor, and don't have any money. Formally, a notion is a subset of the positions in the English string together with the words occupying those positions. To avoid confusion in describing such cases, a subscript will be affixed to each word showing its position. The alignment in FIG. 34 includes the notions the<sub>4</sub> and of<sub>6</sub>the<sub>7</sub>, but not the notions of<sub>6</sub>the<sub>4</sub> or the<sub>7</sub>. In (J'applaudis 'a la décision|I(1) applaud(2) the(4) decision(5)), a is generated by the empty notion. Although the empty notion has no position and so never requires a clarifying subscript, it will be placed at the beginning of the English string, in position zero, and denoted by  $e_0$ . At times, therefore, the previous alignment will be denoted as (J'applaudis à la décision|e<sub>0</sub>(3) I(1) applaud(2) the(4) decision(5)).

The set of all alignments of  $(f|e)$  will be written  $\mathcal{A}(e,f)$ . Since  $e$  has length  $l$  and  $f$  has length  $m$ , there are  $lm$  different connections that can be drawn between them. Since an alignment is determined by the connections that it contains, and since a subset of possible connections can be chosen in  $2^{lm}$  ways, there are  $2^{lm}$  alignments in  $\mathcal{A}(e,f)$ .

For the Models 1-5 described below, only alignments without multi-word notions are allowed. Assuming this restriction, a notion can be identified with a position in the English string, with the empty notion identified with position zero. Thus, each position in the French string is connected to one of the positions 0 through  $l$  in the English string. The notation  $a_j=i$  will be used to indicate that a word in position  $j$  of the French string is connected to the word in position  $i$  of the English string.

The probability of a French string  $f$  and an alignment  $a$  given an English string  $e$  can be written

$$P(f,ale) = P(m|e) \sum_{j=1}^m P(a_j|a_1^{j-1}, f_1^{j-1}, m, e) P(f_j|a_j, f_1^{j-1}, m, e). \quad (47)$$

Here,  $m$  is the length of  $f$  and  $a_1^m$  is determined by  $a$ .

## 9.4 Model 1

The conditional probabilities on the right-hand side of Equation (47) cannot all be taken as independent parameters because there are too many of them. In Model 1, the probabilities  $P(m|e)$  are taken to be independent of  $e$  and  $m$ ; that  $P(a_j|a_1^{j-1}, f_1^{j-1}, m, e)$ , depends only on  $l$ , the length of the English string, and therefore must be  $(l+1)^{-1}$ ; and that  $P(f_j|a_j, f_1^{j-1}, m, e)$  depends only on  $f_j$  and  $e_{a_j}$ . The parameters, then, are  $\epsilon \equiv P(m|e)$ , and  $t(f_j|e_{a_j}) \equiv P(f_j|a_j, f_1^{j-1}, m, e)$ , which will be called the translation probability of  $f_j$  given  $e_{a_j}$ . The parameter  $\epsilon$  is fixed at some small number. The distribution of  $M$  is unnormalized but this is a minor technical issue of no significance to the computations. In particular,  $M$  can be restricted to some finite range. As long as this range encompasses everything that actually occurs in training data, no problems arise.

A method for estimating the translation probabilities for Model 1 will now be described. The joint likelihood of a French sentence and an alignment is

$$P(f,ale) = \epsilon (l+1)^{-m} \prod_{j=1}^m t(f_j|e_{a_j}). \quad (48)$$

The alignment is determined by specifying the values of  $a_j$  for  $j$  from 1 to  $m$ , each of which can take any value from 0 to  $l$ . Therefore,

$$P(fe) = \epsilon (l+1)^{-m} \sum_{a_1=0}^l \dots \sum_{a_m=0}^l \prod_{j=1}^m t(fe_{a_j}). \quad (49)$$

The first goal of the training process is to find the values of the translation probabilities that maximize  $P(fe)$  subject to the constraints that for each  $e$ ,

$$\sum_f t(fe) = 1. \quad (50)$$

An iterative method for doing so will be described.

The method is motivated by the following consideration. Following standard practice for constrained maximization, a necessary relationship for the parameter values at a local maximum can be found by introducing Lagrange multipliers,  $\lambda_e$ , and seeking an unconstrained maximum of the auxiliary function

$$h(t, \lambda) = \epsilon (l+1)^{-m} \sum_{a_1=0}^l \dots \sum_{a_m=0}^l \prod_{j=1}^m t(fe_{a_j}) - \sum_e \lambda_e \left( \sum_f t(fe) - 1 \right). \quad (51)$$

At a local maximum, all of the partial derivatives of  $h$  with respect to the components of  $t$  and  $\lambda$  are zero. That the partial derivatives with respect to the components of  $\lambda$  be zero is simply a restatement of the constraints on the translation probabilities. The partial derivative of  $h$  with respect to  $t(fe)$  is

$$\frac{\partial h}{\partial t(fe)} = \epsilon (l+1)^{-m} \sum_{a_1=0}^l \dots \sum_{a_m=0}^l \sum_{k=1}^m \delta(f, f_j) \delta(e, e_{a_j}) t(fe)^{-1} \prod_{k=1}^m t(fe_{a_k}) - \lambda_e, \quad (52)$$

where  $\delta$  is the Kronecker delta function, equal to one when both of its arguments are the same and equal to zero otherwise. This will be zero provided that

$$t(fe) = \lambda_e^{-1} \epsilon (l+1)^{-m} \sum_{a_1=0}^l \dots \sum_{a_m=0}^l \sum_{k=1}^m \delta(f, f_j) \delta(e, e_{a_j}) \prod_{k=1}^m t(fe_{a_k}). \quad (53)$$

Superficially Equation (53) looks like a solution to the maximization problem, but it is not because the translation probabilities appear on both sides of the equal sign. Nonetheless, it suggests an iterative procedure for finding a solution:

1. Begin with initial guess for the translation probabilities;
2. Evaluate the right-hand side of Equation (53);
3. Use the result as a new estimate for  $t(fe)$ .
4. Iterate steps 2 through 4 until converged.

(Here and elsewhere, the Lagrange multipliers simply serve as a reminder that the translation probabilities must be normalized so that they satisfy Equation (50).) This process, when applied repeatedly is called the EM process. That it converges to a stationary point of  $h$  in situations like this, as demonstrated in the previous section, was first shown by L. E. Baum in an article entitled, An Inequality and Associated

Maximization Technique in Statistical Estimation of Probabilistic Functions of a Markov Process, appearing the journal *Inequalities*, Vol. 3, in 1972.

With the aid of Equation (48), Equation (53) can be re-expressed as

$$t(fe) = \lambda_e^{-1} \sum_a P(f, ae) \underbrace{\sum_{j=1}^m \delta(f, f_j) \delta(e, e_{a_j})}_{\text{number of times } e \text{ connects to } f \text{ in } a}.$$

The expected number of times that  $e$  connects to  $f$  in the translation  $(fe)$  will be called the count of  $f$  given  $e$  for  $(fe)$  and will be denoted by  $c(fe; f, e)$ . By definition,

$$c(fe; f, e) = \sum_a P(ae, f) \sum_{j=1}^m \delta(f, f_j) \delta(e, e_{a_j}), \quad (55)$$

where  $P(ae, f) = P(f, ae) / P(fe)$ . If  $\lambda_e$  is replaced by  $\lambda_e P(fe)$ , then Equation (54) can be written very compactly as

$$t(fe) = \lambda_e^{-1} c(fe; f, e). \quad (56)$$

In practice, the training data consists of a set of translations,  $(f^{(1)}|e^{(1)})$ ,  $(f^{(2)}|e^{(2)})$ ,  $\dots$ ,  $(f^{(S)}|e^{(S)})$ , so this equation becomes

$$t(fe) = \lambda_e^{-1} \sum_{s=1}^S c(fe; f^{(s)}, e^{(s)}). \quad (57)$$

Here,  $\lambda_e$  serves only as a reminder that the translation probabilities must be normalized.

Usually, it is not feasible to evaluate the expectation in Equation (55) exactly. Even if multiword notions are excluded, there are still  $(l+1)^m$  alignments possible for  $(fe)$ . Model 1, however, is special because by recasting Equation (49), it is possible to obtain an expression that can be evaluated efficiently. The right-hand side of Equation (49) is a sum of terms each of which is monomial in the translation probabilities. Each monomial contains  $m$  translation probabilities, one for each of the words in  $f$ . Different monomials correspond to different ways of connecting words in  $f$  to notions in  $e$  with every way appearing exactly once. By direct evaluation, then

$$\sum_{a_1=0}^1 \dots \sum_{a_m=0}^m \pi \tau(f|e_{a_1}) = \pi \sum_{i=0}^1 \tau(f|e_i). \quad (58)$$

Therefore, the sums in Equation (49) can be interchanged with the product to obtain

$$P(f|e) = \epsilon(1+1)^{-m} \pi \sum_{j=1}^m \tau(f|e_j). \quad (59)$$

Using this expression, it follows that

$$c(f|e, f, e) = \frac{\tau(f|e)}{\tau(f|e_0) + \dots + \tau(f|e_l)} \underbrace{\sum_{j=1}^m \tau(f|f_j)}_{\text{count of } f \text{ in } f} \underbrace{\sum_{i=0}^1 \tau(e, e_i)}_{\text{count of } e \text{ in } e}. \quad (60)$$

Thus, the number of operations necessary to calculate a count is proportional to  $lm$  rather than to  $(l+1)^m$  as Equation (55) might suggest.

The details of the initial guesses for  $\tau(f|e)$  are unimportant because  $P(f|e)$  has a unique local maximum for Model 1, as is shown in Section 10. In practice, the initial probabilities  $\tau(f|e)$  are chosen to be equal, but any other choice that avoids zeros would lead to the same final solution.

### 9.5 Model 2

Model 1, takes no cognizance of where words appear in either string. The first word in the French string is just as likely to be connected to a word at the end of the English string as to one at the beginning. In contrast, Model 2 makes the same assumptions as in Model 1 except that it assumes that  $P(a_j|a_1^{j-1}, f_1^{j-1}, m, e)$  depends on  $j$ ,  $a_j$ , and  $m$ , as well as on  $l$ . This is done using a set of alignment probabilities,

$$a(j|j, m, l) \equiv P(a_j|a_1^{j-1}, f_1^{j-1}, m, l), \quad (61)$$

which satisfy the constraints

$$\sum_{i=0}^1 a(i|j, m, l) = 1 \quad (62)$$

for each triple  $jml$ . Equation (49) is replaced by

$$P(f|e) = \epsilon \sum_{a_1=0}^1 \dots \sum_{a_m=0}^m \pi \tau(f|e_{a_1}) a(j|j, m, l). \quad (63)$$

A relationship among the parameter values for Model 2 that maximize this likelihood is obtained, in analogy with the above discussion by seeking an unconstrained maximum of the auxiliary function

$$h(\lambda, \mu) = \epsilon \sum_{a=0}^1 \dots \sum_{a_m=0}^m \pi \tau(f|e_a) a(j|j, m, l) - \sum_e \lambda_e \left( \sum_f \tau(f|e) - 1 \right) - \sum_j \mu_{jml} \left( \sum_i a(i|j, m, l) - 1 \right). \quad (64)$$

It is easily verified that Equations (54), (56), and (57) carry over from Model 1 to Model 2 unchanged. In addition, an iterative update formulas for  $a(i|j, m, l)$ , can be obtained by introducing a new count,  $c(i|j, m, l; f, e)$ , the expected number of times that the word in position  $j$  of  $f$  is connected to the word in position  $i$  of  $e$ . Clearly,

$$c(i|j, m, l; e) = \sum_a P(a|e, f) \delta(i, a_j). \quad (65)$$

In analogy with Equations (56) and (57), for a single translation,

$$a(i|j, m, l) = \mu_{jml}^{-1} c(i|j, m, l; f, e), \quad (66)$$

and, for a set of translations,

$$a(i|j, m, l) = \mu_{jml}^{-1} \sum_{s=1}^S c(i|j, m, l; f^{(s)}, e^{(s)}). \quad (67)$$

Notice that if  $f^{(s)}$  does not have the length  $m$  or if  $e^{(s)}$  does not have length  $l$ , then the corresponding count is zero. As with the  $\lambda$ 's in earlier equations, the  $\mu$ 's here serve to normalize the alignment probabilities.

Model 2 shares with Model 1 the important property that the sums in Equations (55) and (65) can be obtained efficiently. Equation (63) can be rewritten

$$P(f|e) = \epsilon \sum_{j=1}^m \frac{1}{\pi} \tau(f|e_j) a(j|j, m, l). \quad (68)$$

Using this form for  $P(f|e)$ , it follows that

$$c(f|e, f, e) = \sum_{j=1}^m \frac{1}{\pi} \frac{\tau(f|e) a(j|j, m, l) \tau(f|f_j) \tau(e, e_i)}{\tau(f|e_0) a(0|j, m, l) + \dots + \tau(f|e_l) a(l|j, m, l)}, \quad (69)$$

and

$$c(i|j, m, l; f, e) = \frac{\tau(f|e_i) a(i|j, m, l)}{\tau(f|e_0) a(0|j, m, l) + \dots + \tau(f|e_l) a(l|j, m, l)}. \quad (70)$$

Equation (69) has a double sum rather than the product of two single sums, as in Equation (60), because, in Equation (69),  $i$  and  $j$  are tied together through the alignment probabilities.

Model 1 is a special case of Model 2 in which  $a(i|j, m, l)$  is held fixed at  $(l+1)^{-1}$ . Therefore, any set of parameters for Model 1 can be reinterpreted as a set of parameters for Model 2. Taking as initial estimates of the parameters for Model 2 the parameter values that result from training Model 1 is equivalent to computing the probabilities of all alignments using Model 1, but then collecting the counts appropriate to Model 2. The idea of computing the probabilities of the alignments using one model, but collecting the counts in a way appropriate to a second model is very general and can always be used to transfer a set of parameters from one model to another.

### 9.6 Intermodel Interlude

Models 1 and 2 make various approximations to the conditional probabilities that appear in Equation (47). Although Equation (47) is an exact statement, it is only one of many ways in which the joint likelihood of  $f$  and  $e$  can be written as a product of conditional probabilities. Each such product corresponds in a natural way to a generative process for developing  $f$  and  $e$ . In the process corresponding

to Equation (47), a length for  $f$  is chosen first. Next, a position in  $e$  is selected and connected to the first position in  $f$ . Next, the identity of  $f_1$  is selected. Next, another position in  $e$  is selected and this is connected to the second word in  $f$ , and so on.

Casual inspection of some translations quickly establishes that the is usually translated into a single word (le, la or l'), but is sometimes omitted; or that only is often translated into one word (for example, seulement), but sometimes into two (for example, ne . . . que), and sometimes into none. The number of French words to which  $e$  is connected in an alignment is a random variable,  $\Phi_e$ , that will be called its fertility. Each choice of the parameters in Model 1 or Model 2 determines a distribution,  $\Pr(\Phi_e = \phi)$ , for this random variable. But the relationship is remote: just what change will be wrought in the distribution of  $\Phi_{the}$  if, say,  $a(112, 8, 9)$  is adjusted, is not immediately clear. In Models 3, 4, and 5, fertilities are parameterized directly.

As a prolegomenon to a detailed discussion of Models 3, 4, and 5, the generative process upon which they are based will be described. Given an English string,  $e$ , a fertility for each word is selected, and a list of French words to connect to it is generated. This list, which may be empty, will be called a tablet. The collection of tablets is a random variable,  $T$ , that will be called the tableau of  $e$ ; the tablet for the  $i^{th}$  English word is a random variable,  $T_i$ ; and the  $k^{th}$  French word in the  $i^{th}$  tablet is a random variable,  $T_{ik}$ . After choosing the tableau, a permutation of its words is generated, and  $f$  is produced. This permutation is a random variable,  $\Pi$ . The position of  $f$  of the  $k^{th}$  word in the  $i^{th}$  tablet is yet another a random variable,  $\Pi_{ik}$ .

The joint likelihood for a tableau,  $\tau$ , and a permutation,  $\pi$ , is

$$P(\tau, \pi | e) = \prod_{i=1}^l P(\phi_i | \phi_1^{i-1}, e) P(\phi_0 | \phi_1^l, e) \times \prod_{i=0}^l \prod_{k=1}^{\phi_i} P(\tau_{ik} | \tau_{i1}^{k-1}, \tau_0^{i-1}, \phi_0^l, e) \times \prod_{i=1}^l \prod_{k=1}^{\phi_i} P(\pi_{ik} | \pi_{i1}^{k-1}, \pi_1^{i-1}, \tau_0^l, \phi_0^l, e) \times \prod_{k=1}^{\phi_0} P(\pi_{0k} | \pi_{01}^{k-1}, \pi_1^l, \tau_0^l, \phi_0^l, e). \quad (71)$$

In this equation,  $\tau_{i1}^{k-1}$  represents the series of values  $\tau_{i1}, \dots, \tau_{ik-1}$ ;  $\pi_{i1}^{k-1}$  represents the series of values  $\pi_{i1}, \dots, \pi_{ik-1}$ ; and  $\phi_{e1}$ .

Knowing  $\tau$  and  $\pi$  determines a French string and an alignment, but in general several different pairs  $\tau, \pi$  may lead to the same pair  $f, a$ . The set of such pairs will be denoted by  $(f, a)$ . Clearly, then

$$P(f, a | e) = \sum_{(\tau, \pi) \in \langle f, a \rangle} P(\tau, \pi | e). \quad (72)$$

The number of elements in  $(f, a)$  is

$$\prod_{i=0}^l \phi_i!$$

because for each  $\tau_i$  there are  $\phi_i!$  arrangements that lead to the pair  $f, a$ . FIG. 36 shows the two tableaux for (bon marché)cheap(1,2)).

Except for degenerate cases, there is one alignment in  $\mathcal{A}(e, f)$  for which  $P(a | e, f)$  is greatest. This alignment will be called the Viterbi alignment for  $(f, e)$  and will be denoted by  $V(f, e)$ . For Model 2 (and, thus, also for Model 1), finding  $V(f, e)$  is straightforward. For each  $j$ ,  $a_j$  is chosen so as to make the product  $t(f_j | e_{a_j}) a(a_j | j, m, l)$  as large as possible. The Viterbi alignment depends on the model with respect to which it is computed. In what follows, the Viterbi alignments for the different models will be distinguished by writing  $V(f, e; 1)$ ,  $V(f, e; 2)$ , and so on.

The set of all alignment for which  $a_j = i$  will be denoted by  $\mathcal{A}_i \leftarrow j(e, f)$ . The pair  $ij$  will be said to be pegged for the elements of  $\mathcal{A}_i \leftarrow j(e, f)$ . The element of  $\mathcal{A}_i \leftarrow j(e, f)$  for which  $P(a | e, f)$  is greatest will be called the pegged Viterbi alignment for  $ij$ , and will be denoted by  $V_{i \leftarrow j}(f, e)$ . Obviously,  $V_{i \leftarrow j}(f, e; 1)$  and  $V_{i \leftarrow j}(f, e; 2)$  can be found quickly with a straightforward modification of the method described above for finding  $V(f, e; 1)$  and  $V(f, e; 2)$ .

### 9.7 Model 3

Model 3 is based on Equation (71). It assumes that, for  $i$  between 1 and  $l$ ,  $P(\phi_i | \phi_1^{i-1}, e)$  depends only on  $\phi_i$  and  $e_i$ ; that, for all  $i$ ,  $P(\tau_{ik} | \tau_{i1}^{k-1}, \tau_0^{i-1}, \phi_0^l, e)$  depends only on  $\tau_{ik}$  and  $e_i$ ; and that, for  $i$  between 1 and  $l$ ,  $P(\pi_{ik} | \pi_{i1}^{k-1}, \pi_1^{i-1}, \tau_0^l, \phi_0^l, e)$  depends only on  $\pi_{ik}$ ,  $i$ ,  $m$ , and  $l$ . The parameters of Model 3 are thus a set of fertility probabilities,  $n(\phi_i | e) = P(\phi_i | \phi_1^{i-1}, e)$ ; a set of translation probabilities,  $t(f, e) = \Pr(T_{ik} = f | \tau_{i1}^{k-1}, \tau_0^{i-1}, \phi_0^l, e)$ ; and a set of distortion probabilities,  $d(j | i, m, l) = \Pr(\Pi_{ik} = j | \pi_{i1}^{k-1}, \pi_1^{i-1}, \tau_0^l, \phi_0^l, e)$ .

The distortion and fertility probabilities for  $e_0$  are handled differently. The empty notion conventionally occupies position 0, but actually has no position. Its purpose is to account for those words in the French string that cannot readily be accounted for by other notions in the English string. Because these words are usually spread uniformly throughout the French string, and because they are placed only after all of the other words in the sentence have been placed, the probability  $\Pr(\Pi_{0k+1} = j | \pi_{01}^k, \pi_1^l, \tau_0^l, \phi_0^l, e)$  is set to 0 unless position  $j$  is vacant in which case it is set  $(\phi_0 - k)^{-1}$ . Therefore, the contribution of the distortion probabilities for all of the words in  $\tau_0$  is  $1/\phi_0!$ .

The value of  $\phi_0$  depends on the length of the French sentence since longer sentences have more of these extraneous words. In fact Model 3 assumes that

$$P(\phi_0 | \phi_1^l, e) = \binom{\phi_1 + \dots + \phi_l}{\phi_0} p_0^{\phi_1 + \dots + \phi_l - \phi_0} p_1^{\phi_0} \quad (73)$$

for some pair of auxiliary parameters  $p_0$  and  $p_1$ . The expression on the left-hand side of this equation depends on  $\phi_1^l$  only through the sum  $\phi_1 + \dots + \phi_l$  and defines a probability distribution over  $\phi_0$  whenever  $p_0$  and  $p_1$  are nonnegative and sum to 1. The probability  $P(\phi_0 | \phi_1^l, e)$  can be interpreted as follows. Each of the words from  $\tau_1^l$  is imagined to require an extraneous word with probability  $p_1$ ; this word is required to be connected to the empty notion. The probability that



exactly  $\phi_0$  of the words from  $\tau_1$  will require an extraneous word is just the expression given in Equation (73).

As in Models 1 and 2, an alignment of (fle) in Model 3 is determined by specifying  $a_j$  for each position in the French string. The fertilities,  $\phi_0$  through  $\phi_i$ , are functions of the  $a_j$ 's. Thus,  $\phi_i$  is equal to the number of  $j$ 's for which  $a_j$  equals  $i$ . Therefore,

$$P(fle) = \sum_{a_1=0}^1 \dots \sum_{a_m=0}^1 P(f, ale) \quad (74)$$

$$= \sum_{a_1=0}^1 \dots \sum_{a_m=0}^1 \binom{m-\phi_0}{\phi_0} p_0^{m-2\phi_0} p_1^{\phi_0} \prod_{i=1}^i \phi_i! n(\phi_i | e_i) \times \prod_{j=1}^m \tau(f | e_{a_j}) d(j | a_j, m, l)$$

with  $\sum_i \tau(f | e_i) = 1$ ,  $\sum_i d(j | i, m, l) = 1$ ,  $\sum_i n(\phi_i | e_i) = 1$ , and  $p_0 + p_1 = 1$ . According to the assumptions of Model 3, each of the pairs  $(\tau, \pi)$  in  $(f, a)$  makes an identical contribution to the sum in Equation (72). The factorials in Equation (74) come from carrying out this sum explicitly. There is no factorial for the empty notion because it is exactly cancelled by the contribution from the distortion probabilities.

By now, it should be clear how to provide an appropriate auxiliary function for seeking a constrained maximum of the likelihood of a translation with Model 3. A suitable auxiliary function is

$$h(i, d, n, p, \lambda, \mu, v, \xi) = \quad (75)$$

$$P(f | e) - \sum_e \lambda_e \left( \sum_f \tau(f | e) - 1 \right) - \sum_i \mu_{mi} (\sum_j d(j | i, m, l) - 1) -$$

$$\sum_e v_e \left( \sum_{\phi} n(\phi | e) - 1 \right) - \xi(p_0 + p_1 - 1). \quad (76)$$

As with Models 1 and 2, counts are defined by

$$c(f | e; f, e) = \sum_a P(a | e, f) \sum_{j=1}^m \delta(f | f_j) \delta(e | e_{a_j}), \quad (76)$$

$$c(j | i, m, l; f, e) = \sum_a P(a | e, f) \delta(i, a_j), \quad (77)$$

$$c(\phi | e; f, e) = \sum_a P(a | e, f) \sum_{i=1}^i \delta(\phi, \phi_i) \delta(e | e_i), \quad (78)$$

$$c(0; f, e) = \sum_a P(a | e, f) (m - 2\phi_0) \quad (79)$$

and

$$c(1; f, e) = \sum_a P(a | e, f) \phi_0. \quad (80)$$

In terms of these counts, the reestimation formulae for Model 3 are

$$\tau(f | e) = \lambda_e^{-1} \sum_{s=1}^S c(f | e; f^{(s)}, e^{(s)}), \quad (81)$$

$$d(j | i, m, l) = \mu_{mi}^{-1} \sum_{s=1}^S c(j | i, m, l; f^{(s)}, e^{(s)}), \quad (82)$$

$$n(\phi | e) = v_e^{-1} \sum_{s=1}^S c(\phi | e; f^{(s)}, e^{(s)}), \quad (83)$$

and

-continued

$$p_i = \xi^{-1} \sum_{s=1}^S c(k; f^{(s)}, e^{(s)}). \quad (84)$$

Equations (76) and (81) are identical to Equations (55) and (57) and are repeated here only for convenience. Equations (77) and (82) are similar to Equations (65) and (67), but  $a(i | j, m, l)$  differs from  $d(j | i, m, l)$  in that the former sums to unity over all  $i$  for fixed  $j$  while the latter sums to unity over all  $j$  for fixed  $i$ . Equations (78), (79), (80), (83), and (84), for the fertility parameters, are new.

The trick that permits the a rapid evaluation of the right-hand sides of Equations (55) and (65) efficiently for Model 2 does not work for Model 3. Because of the fertility parameters, it is not possible to exchange the sums over  $a_1$  through  $a_m$  with the product over  $j$  in Equation (74) as was done for Equations (49) and (63). A method for overcoming this difficulty now will be described. This method relies on the fact that some alignments are much more probable than others. The method is to carry out the sums in Equations (74) and (76) through (80) only over some of the more probable alignments, ignoring the vast sea of much less probable ones.

To define the subset,  $S$ , of the elements of  $\mathcal{A}(fle)$  over which the sums are evaluated a little more notation is required. Two alignments,  $a$  and  $a'$  will be said to differ by a move if there is exactly one value of  $j$  for which  $a_j \neq a'_j$ . Alignments will be said to differ by a swap if  $a_j = a'_j$  except at two values,  $j_1$  and  $j_2$ , for which  $a_{j_1} = a'_{j_2}$  and  $a_{j_2} = a'_{j_1}$ . The two alignments will be said to be neighbors if they are identical or differ by a move or by a swap. The set of all neighbors of  $a$  will be denoted by  $\mathcal{N}(a)$ .

Let  $b(a)$  be that neighbor of  $a$  for which the likelihood is greatest. Suppose that  $ij$  is pegged for  $a$ . Among the neighbors of  $a$  for which  $ij$  is also pegged, let  $b_{i \leftarrow j}(a)$  be that for which the likelihood is greatest. The sequence of alignments  $a, b(a), b^2(a) \equiv b(b(a)), \dots$ , converges in a finite number of steps to an alignment that will be denoted as  $b^{**}(a)$ . Similarly, if  $ij$  is pegged for  $a$ , the sequence of alignments  $a, b_{i \leftarrow j}(a), b_{i \leftarrow j}^2(a), \dots$ , converges in a finite number of steps to an alignment that will be denoted as  $b_{i \leftarrow j}^{**}(a)$ . The simple form of the distortion probabilities in Model 3 makes it easy to find  $b(a)$  and  $b_{i \leftarrow j}(a)$ . If  $a'$  is a neighbor of  $a$  obtained from it by the move of  $j$  from  $i$  to  $i'$ , and if neither  $i$  nor  $i'$  is 0, then

$$P(a'e,f) = P(a'e,f) \frac{(\phi_i + 1)}{\phi_i} \frac{n(\phi_i + 1e_f)}{n(\phi_i e_f)} \frac{n(\phi_i - 1e_d)}{n(\phi_i e_d)} \frac{t(f|e_f)}{t(f|e_d)} \frac{d(j|i', m, l)}{d(j|i, m, l)} \quad (85)$$

Notice that  $\phi_i$  is the fertility of the word in position  $i'$  for alignment  $a$ . The fertility of this word in alignment  $a'$  is  $\phi_i + 1$ . Similar equations can be easily derived with either  $i$  or  $i'$  is zero, or when  $a$  and  $a'$  differ by a swap.

Given these preliminaries,  $S$  is defined by

$$= (b^{**}(V(f|e;2))) \cup_{ij} (b_{i \leftarrow j}^{**}(V_{i \leftarrow j}(f|e;2))). \quad (86)$$

In this equation,  $b^{**}(V(f|e;2))$  and  $b_{i \leftarrow j}^{**}(V_{i \leftarrow j}(f|e;2))$  are approximations to  $V(f|e;3)$  and  $V_{i \leftarrow j}(f|e;3)$  neither of which can be computed efficiently.

In one iteration of the EM process for Model 3, the counts in Equations (76) through (80), are computed by summing only over elements of  $S$ . These counts are then used in Equations (81) through (84) to obtain a new set of parameters. If the error made by including only some of the elements of  $A(e,f)$  is not too great, this iteration will lead to values of the parameters for which the likelihood of the training data is at least as large as for the first set of parameters.

The initial estimates of the parameters of Model 2 are adapted from the final iteration of the EM process for Model 2. That is, the counts in Equations (76) through (80) are computed using Model 2 to evaluate  $P(a|e,f)$ . The simple for of Model 2 again makes exact calculation feasible. The Equations (69) and (70) are readily adapted to compute counts for the translation and distortion probabilities; efficient calculation of the fertility counts is more involved. A discussion of how this is done is given in Section 10.

### 9.8 Deficiency

A problem with the parameterization of the distortion probabilities in Model 3 is this: whereas the sum over all pairs  $\tau, \pi$  of the expression on the right-hand side of Equation (71) is unity, if  $\Pr(\Pi_{i=1}^k = j | \pi_{i=1}^{k-1}, \pi_1^{i-1}, \tau_0^i, \phi_0^i, e)$  depends only on  $j, i, m$ , and  $l$  for  $i > 0$ . Because the distortion probabilities for assigning positions to later words do not depend on the positions assigned to earlier words, Model 3 wastes some of its probability on what will be called generalized strings, i.e., strings that have some positions with several words and others with none. When a model has this property of not concentrating all of its probability on events of interest, it will be said to be deficient. Deficiency is the price for the simplicity that permits Equation (85).

Deficiency poses no serious problem here. Although Models 1 and 2 are not technically deficient, they are surely spiritually deficient. Each assigns the same probability to the alignments (Je n'ai pas de stylo | I do not(2,4) have(3) a(5) pen(6)) and (Je pas ai ne de stylo | I do not(2,4) have(3) a(5) pen(6)), and, therefore, essentially the same probability to the translations (Je n'ai pas de stylo | I do not have a pen) and (Je pas ai ne de stylo | I do not have a pen). In each case, not produces two words, ne and pas, and in each case, one of these words ends up in the second position of the French string and the other in the fourth position. The first translation should be much more probable than the second, but this defect is of little concern because while the system may be required to translate the first string someday, it will almost surely not be required to translate the second. The translation

models are not used to predict French given English but rather as a component of a system designed to predict English given French. They need only be accurate to within a constant factor over well-formed strings of French words.

### 9.9 Model 4

Often the words in an English string constitute phrases that are translated as units into French. Sometimes, a translated phrase may appear at a spot in the French string different from that at which the corresponding English phrase appears in the English string. The distortion probabilities of Model 3 do not account well for this tendency of phrases to move around as units. Movement of a long phrase will be much less likely than movement of a short phrase because each word must be moved independently. In Model 4, the treatment of  $\Pr(\Pi_{i=1}^k = j | \pi_{i=1}^{k-1}, \pi_1^{i-1}, \tau_0^i, \phi_0^i, e)$  is modified so as to alleviate this problem. Words that are connected to the empty notion do not usually form phrases and so Model 4 continues to assume that these words are spread uniformly throughout the French string.

As has been described, an alignment resolves an English string into a notional scheme consisting of a set of possibly overlapping notions. Each of these notions then accounts for one or more French words. In Model 3 the notional scheme for an alignment is determined by the fertilities of the words: a word is a notion of its fertility is greater than zero. The empty notion is a part of the notional scheme if  $\phi_0$  is greater than zero. Multi-word notions are excluded. Among the one-word notions, there is a natural order corresponding to the order in which they appear in the English string. Let  $[i]$  denote the position in the English string of the  $i^{\text{th}}$  one-word notion. The center of this notion,  $\odot_i$ , is defined to be the ceiling of the average value of the positions in the French string of the words from its tablet. The head of the notion is defined to be that word in its tablet for which the position in the French string is smallest.

In Model 4, the probabilities  $d(j|i, m, l)$  are replaced by two sets of parameters: one for placing the head of each notion, and one for placing any remaining words. For  $[i] > 0$ , Model 4 requires that the head for notion  $i$  be  $\tau_{[i]1}$ . It assumes that

$$\Pr(\Pi_{[i]1} = j | \pi_{[i]1}^{[i]-1}, \tau_0^i, \phi_0^i, e) = d_1(j - \odot_{[i]-1}, A(e_{[i]-1}), B(f_j)). \quad (87)$$

Here,  $A$  and  $B$  are functions of the English and French word that take on a small number of different values as their arguments range over their respective vocabularies. In the Section entitled Classes, a process is described for dividing a vocabulary into classes so as to preserve mutual information between adjacent classes in running text. The classes

$A$  and  $B$  are constructed as functions with fifty distinct values by dividing the English and French vocabularies each into fifty classes according to this method. The probability is assumed to depend on the previous notion and on the identity of the French word being placed so as to account for such facts as the appearance of adjectives before nouns in English but after them in French. The displacement for the head of notion  $i$  is denoted by  $j - \odot_{[i]-1}$ . It may be either positive or negative. The probability of  $d_1(-1 | A(e) B(f))$  is expected to be larger than  $d_1(+1 | A(e) B(f))$  when  $e$  is an adjective and  $f$  is a noun. Indeed, this is borne out in the

trained distortion probabilities for Model 4, where  $d_1(-1)A(\text{government's}) B(\text{développement})$  is 0.9269, while  $d_1(+1)A(\text{government's}) B(\text{développement})$  is 0.0022.

The placement of the  $k^{\text{th}}$  word of notion  $i$  for  $[i]>0$ ,  $k>1$  is done as follows. It is assumed that

$$Pr(I_{[i]k}=j | \pi_{[i]k-1}^{(i-1)}, \pi_{[i]k-1}^{(i-1)}, \tau_{[i]k-1}^{(i-1)}, \tau_{[i]k-1}^{(i-1)}, e) = d_{-1}(j - \pi_{[i]k-1}^{(i-1)} | B(j)). \quad (88)$$

The position  $\pi_{[i]k}$  is required to be greater than  $\pi_{[i]k-1}$ . Some English words tend to produce a series of French words that belong together, while others tend to produce a series of words that should be separate. For example, implemented can produce mis en application, which usually occurs as a unit, but not can produce ne pas, which often occurs with an intervening verb. Thus  $d_{-1}(2|B(\text{pas}))$  is expected to be relatively large compared with  $d_{-1}(2|B(\text{en}))$ . After training, it is in fact true that  $d_{-1}(2|B(\text{pas}))$  is 0.6936 and  $d_{-1}(2|B(\text{en}))$  is 0.0522.

It is allowed to place  $\tau_{[i]k}$  either before or after any previously positioned words, but subsequent words from  $\tau_{[i]k}$  are required to be placed in order. This does not mean that they must occupy consecutive positions but only that the second word from  $\tau_{[i]k}$  must lie to the right of the first, the third to the right of the second, and so on. Because of this, at most one of the  $\phi_{[i]k}!$  arrangements of  $\tau_{[i]k}$  is possible.

The count and reestimation formulae for Model 4 are similar to those for the previous Models and will not be given explicitly here. The general formulae in Section 10 are helpful in deriving these formulae. Once again, the several counts for a translation are expectations of various quantities over the possible alignments with the probability of each alignment computed from an earlier estimate of the parameters. As with Model 3, these expectations are computed by sampling some small set,  $S$ , of alignments. As described above, the simple form that for the distortion probabilities in Model 3, makes it possible to find  $b^m(a)$  rapidly for any  $a$ . The analogue of Equation (85) for Model 4 is complicated by the fact that when a French word is moved from notion to notion, the centers of two notions change, and the contribution of several words is affected. It is nonetheless possible to evaluate the adjusted likelihood incrementally, although it is substantially more time consuming.

Faced with this unpleasant situation, the following method is employed. Let the neighbors of  $a$  be ranked so that the first is the neighbor for which  $P(\text{ale}, f; 3)$  is greatest, the second the one for which  $P(\text{ale}, f; 3)$  is next greatest, and so on. Then  $\tilde{b}(a)$  is the highest ranking neighbor of  $a$  for which  $P(\tilde{b}(a)|e, f; 4)$  is at least as large as  $P(\text{alf}, e; 4)$ . Define  $\tilde{b}_{i \leftarrow j}(a)$  analogously. Here,  $P(\text{ale}, f; 3)$  means  $P(\text{ale}, f)$  as computed with Model 3, and  $P(\text{ale}, f; 4)$  means  $P(\text{ale}, f)$  as computed with Model 4. Define  $S$  for Model 4 by

$$= (b^m(V(f; 2))) \cup \bigcup_{ij} (\tilde{b}_{i \leftarrow j}(V(f; 2))). \quad (89)$$

This equation is identical to Equation (89) except that  $b$  has been replaced by  $\tilde{b}$ .

#### 9.10 Model 5

Like Model 3 before it, Model 4 is deficient. Not only can several words lie on top of one another, but words can be placed before the first position or beyond the last position in the French string. Model 5 removes this deficiency.

After the words for  $\tau_{[i]k-1}$  and  $\tau_{[i]k-1}^{k-1}$  are placed, there will remain some vacant positions in the French string. Obviously,  $\tau_{[i]k}$  should be placed in one of these vacancies. Models 3 and 4 are deficient precisely because they fail to enforce this constraint for the one-word notions. Let  $v(j, \tau_{[i]k-1}^{(i-1)}, \tau_{[i]k-1}^{k-1})$  be the number of vacancies up to and including position  $j$  just before  $\tau_{[i]k}$  is placed. This position will be denoted simply as  $v_j$ . Two sets of distortion parameters are retained as in Model 4; they will again be referred to as  $d_1$  and  $d_{-1}$ . For  $[i]>0$ , it is required that

$$Pr(I_{[i]k}=j | \pi_{[i]k-1}^{(i-1)}, \tau_{[i]k-1}^{(i-1)}, \tau_{[i]k-1}^{k-1}, e) = d_1(v_j | B(f_j, v_{[i]k-1}, v_m \times (1 - \delta(v_j, v_{[i]k-1}))). \quad (90)$$

The number of vacancies up to  $j$  is the same as the number of vacancies up to  $j-1$  only when  $j$  is not, itself, vacant. The last factor, therefore, is 1 when  $j$  is vacant and 0 otherwise. As with Model 4,  $d_1$  is allowed to depend on the center of the previous notion and on  $f_j$ , but is not allowed to depend on  $e_{[i]k-1}$  since doing so would result in too many parameters. For  $[i]>0$  and  $k>1$ , it is assumed that

$$Pr(I_{[i]k}=j | \pi_{[i]k-1}^{(i-1)}, \pi_{[i]k-1}^{(i-1)}, \tau_{[i]k-1}^{(i-1)}, \tau_{[i]k-1}^{k-1}, e) = d_{-1}(v_j - v_{\pi_{[i]k-1}^{(i-1)}} | B(f_j, v_m - v_{\pi_{[i]k-1}^{(i-1)}})). \quad (91)$$

Again, the final factor enforces the constraint that  $\tau_{[i]k}$  land in a vacant position, and, again, it is assumed that the probability depends on  $f_j$  only through its class.

The details of the count and reestimation formulae are similar to those for Model 4 and will not be explained explicitly here. Section 10 contains useful equations for deriving these formulae. No incremental evaluation of the likelihood of neighbors is possible with Model 5 because a move or swap may require wholesale recomputation of the likelihood of an alignment. Therefore, when the expectations for Model 5 are evaluated, only the alignments in  $S$  as defined in Equation (89) are included. The set of alignments included in the sum is further trimmed by removing any alignment,  $a$ , for which  $P(\text{ale}, f; 4)$  is too much smaller than  $P(b^m(V(f; 2)|e, f; 4))$ .

Model 5 provides a powerful tool for aligning translations. Its parameters are reliably estimated by making use of Models 2, 3 and 4. In fact, this is the *raison d'être* of these models. To keep them aware of the lay of the land, their parameters are adjusted as the iterations of the EM process for Model 5 are performed. That is, counts for Models 2, 3, and 4 are computed by summing over alignments as determined by the abbreviated  $S$  described above, using Model 5 to compute  $P(\text{ale}, f)$ . Although this appears to increase the storage necessary for maintaining counts as the training is processed, the extra burden is small because the overwhelming majority of the storage is devoted to counts for  $t(f|e)$ , and these are the same for Models 2, 3, 4, and 5.

#### 9.11 Results

A large collection of training data is used to estimate the parameters of the five models described above. In one embodiment of these models, training data is obtained using the method described in detail in the paper, *Aligning Sentences in Parallel Corpora*, by P. F. Brown, J. C. Lai, and R. L. Mercer, appearing in the *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics*, June 1991. This paper is incorporated by reference herein. This method is applied to a large number of translations from several years of the proceedings of the Canadian parliament. From these translations, a training data set is chosen comprising those pairs for which both the English sentence and the French sentence are thirty words or less in length. This is a collection of 1,778,620 translations. In an

effort to eliminate some of the typographical errors that abound in the text, a English vocabulary is chosen consisting of all of those words that appear at least twice in English sentences in the data, and as a French vocabulary is chosen consisting of all those words that appear at least twice in French sentences in the data. All other words are replaced with a special unknown English word or unknown French word according as they appear in an English sentence or a French sentence. In this way an English vocabulary of 42,005 words and a French vocabulary of 58,016 words is obtained. Some typographic errors are quite frequency, for example, momento for memento, and so the vocabularies are not completely free of them. At the same time, some words are truly rare, and in some cases, legitimate words are omitted. Adding  $e_0$  to the English vocabulary brings it to 42,006 words.

Eleven iterations of the EM process are performed for this data. The process is initialized by setting each of the 2,437,020,096 translation probabilities,  $t(f|e)$ , to  $1/58016$ . That is, each of the 58,016 words in the French vocabulary is assumed to be equally likely as a translation for each of the 42,006 words in the English vocabulary. For  $t(f|e)$  to be greater than zero at the maximum likelihood solution for one of the models,  $f$  and  $e$  must occur together in at least one of the translations in the training data. This is the case for only 25,427,016 pairs, or about one percent of all translation probabilities. On the average, then, each English word appears with about 605 French words.

Table 6 summarizes the training computation. At each iteration, the probabilities of the various alignments of each translation using one model are computed, and the counts using a second, possibly different model are accumulated. These are referred to in the table as the In model and the Out model, respectively. After each iteration, individual values are retrained for only those translation probabilities that surpass a threshold; the remainder are set to the small value ( $10^{-12}$ ). This value is so small that it does not affect the normalization conditions, but is large enough that translation probabilities can be resurrected during later

TABLE 6

A summary of the training iterations.				
Iteration	In → Out	Survivors	Alignments	Perplexity
	2 → 1	12,017,609		71550.56
2	2 → 2	12,160,475		202.99
3	2 → 2	9,403,220		89.41
4	2 → 2	6,837,172		61.59
5	2 → 2	5,303,312		49.77
6	2 → 2	4,397,172		46.36
7	2 → 3	3,841,470		45.15
8	3 → 5	2,057,033	291	124.28
9	5 → 5	1,850,665	95	39.17
10	5 → 5	1,763,665	48	32.91
11	5 → 5	1,703,393	39	31.29
12	5 → 5	1,658,364	33	30.65

iterations. As is apparent from columns 4 and 5, even though the threshold is lowered as iterations progress, fewer and fewer probabilities survive. By the final iteration, only 1,620,287 probabilities survive, an average of about thirty-nine French words for each English word.

As has been described, when the In model is neither Model 1 nor Model 2, the counts are computed by summing over only some of the possible alignments. Many of these alignments have a probability much smaller than that of the Viterbi alignment. The column headed Alignments in Table 6 shows the average number of alignments for which the

probability is within a factor of 25 of the probability of the Viterbi alignment in each iteration. As this number drops, the model concentrates more and more probability onto fewer and fewer alignments so that the Viterbi alignment becomes ever more dominant.

The last column in the table shows the perplexity of the French text given the English text for the In model of the iteration. The likelihood of the training data is expected to increase with each iteration. This likelihood can be thought of as arising from a product of factors, one for each French word in the training data. There are 28,850,104 French words in the training data so the 28,850,104<sup>th</sup> root of the likelihood is the average factor by which the likelihood is reduced for each additional French word. The reciprocal of this root is the perplexity shown in the table. As the likelihood increases, the perplexity decreases. A steady decrease in perplexity is observed as the iterations progress except when a switch from Model 2 as the In model to Model 3 is made. This sudden jump is not because Model 3 is a poorer model than Model 2, but because Model 3 is deficient; the great majority of its probability is squandered on objects that are not strings of French words. As has been explained, deficiency is not a problem. In the description of Model 1, the  $P(m|e)$  was left unspecified. In quoting perplexities for Models 1 and 2, it is assumed that the length of the French string is Poisson with a mean that is a linear function of the length of the English string. Specifically, it is assumed that  $\Pr(M=m|e) = (\lambda l)^m e^{-\lambda l} / m!$ , with  $\lambda$  equal to 1.09.

It is interesting to see how the Viterbi alignments change as the iterations progress. In Table 7, the Viterbi alignment for several sentences after iterations 1, 6, 7, and 11 are displayed. These are the final iterations for Models 1, 2, 3, and 5, respectively. In each example, a subscripted *in* is affixed to each word to help in interpreting the list of numbers after each English word. In the first example, (*Il me semble faire signe que ouillt* seems to me that he is nodding), two interesting changes evolve over the course of the iterations. In the alignment for Model 1, *Il* is correctly connected to *he*, but in all later alignments *Il* is incorrectly connected to *It*. Models 2, 3, and 5 discount a connection of *he* to *Il* because it is quite far away. None of the five models is sophisticated enough to make this connection properly. On the other hand, while nodding, is connected only to *signe* by Models 1, 2, and 3, and *oui*, it is correctly connected to the entire phrase *faire signe que oui* by Model 5. In the second example, (*Voyez les profits que ils ont réalisés*! Look at the profits they have made), Models 1, 2, and 3 incorrectly connect *profits*, to both *profits*, and *réalisés*, but with Model 5, *profits*, is correctly connected to *profits*, and *made*, is connected to *réalisés*,. Finally, in (*De les promesses, de les promesses!* Promises, promises.), *Promises*, is connected to both instances of *promesses* with Model 1; *promises*, is connected to most of the French sentence with Model 2; the final punctuation of the English sentence is connected to both the exclamation point and, curiously, to *de*, with Model 3. Only Model 5 produces a satisfying alignment of the two sentences. The orthography for the French sentence in the second example is *Voyez les profits qu'ils ont réalisés* and in the third example is *Des Promesses, des promesses!* Notice that the *e* has been restored to the end of *qu'* and *des* has twice been analyzed into its constituents, *de* and *les*. These and other petty pseudographic improprieties are committed in the interest of regularizing the French text. In all cases, orthographic French can be recovered by rule from the corrupted versions.

Tables 8 through 17 show the translation probabilities and fertilities after the final iteration of training for a number of English words. All and only those probabilities that are greater than 0.01 are shown. Some words, like nodding, in Table 8, do not slip gracefully into French. Thus, there are translations like (Il fait signe que oui) He is nodding), (Il fait un signe de la tête) He is nodding), (Il fait un signe de tête affirmatif) He is nodding), or (Il hoche la tête affirmativement) He is nodding). As a result, nodding frequently has a large fertility and spreads its translation probability over a variety of words. In French, what is worth saying is worth saying in many different ways. This is also seen with words like should, in Table 9, which rarely has a fertility greater than one but still produces many different words, among them devrait, devraient, devrions, doit, doivent, devons, and devrais. These are (just a fraction of the many) forms of the French verb devoir. Adjectives fare a little better: national, in Table 10, almost never produces more than one word and confines itself to one of nationale, national, nationaux, and nationales, respectively the feminine, the masculine, the masculine plural, and the feminine plural of the corresponding French adjective. It is clear that the models would benefit from some kind of morphological processing to rein in the lexical exuberance of French.

As seen from Table 11, the produces le, la, les, and l' as is expected. Its fertility is usually 1, but in some situations English prefers an article where French does not and so about 14% of the time its fertility is 0. Sometimes, as with farmers, in Table 12, it is French that prefers the article. When this happens, the English noun trains to product its translation

TABLE 7

The progress of alignments with iteration.			
Il, me <sub>2</sub> semble <sub>3</sub> faire <sub>4</sub> signe <sub>5</sub> que <sub>6</sub> oui <sub>7</sub>			
It seems(3) to(4) me(2) that(6) he(1) is nodding(5,7)			
It(1) seems(3) to me(2) that he is nodding(5,7)			
It(1) seems(3) to(4) me(2) that(6) he is nodding(5,7)			
It(1) seems(3) to me(2) that he is nodding(4,5,6,7)			
Voyez, les <sub>2</sub> profits <sub>3</sub> que <sub>4</sub> ils <sub>5</sub> ont <sub>6</sub> réalisés <sub>7</sub>			
Look(1) at the(2) profits(3,7) they(5) have(6) made			
Look(1) at the(2,4) profits(3,7) they(5) have(6) made			
Look(1) at the profits(3,7) they(5) have(6) made			
Look(1) at the(2) profits(3) they(5) have(6) made(7)			
De, les <sub>2</sub> promesses <sub>3</sub> , 4 de, les <sub>5</sub> promesses <sub>6</sub> , l' <sub>8</sub>			
Promises(3,7) ,(4) promises ,(8)			
Promises(4) promises(2,3,6,7) ,(8)			
Promises(3) ,(4) promises(7) ,(5,8)			
Promises(2,3) ,(4) promises(6,7) ,(8)			

TABLE 8

Translation and fertility probabilities for nodding. nodding			
f	t(f e)	φ	n(φ e)
signe	0.164	4	0.342
la	0.123	3	0.293
tête	0.097	2	0.167
oui	0.086	1	0.163
fait	0.073	0	0.023
que	0.073		
hoche	0.054		
hocher	0.048		
faire	0.030		

TABLE 8-continued

Translation and fertility probabilities for nodding. nodding			
f	t(f e)	φ	n(φ e)
me	0.024		
approuve	0.019		
qui	0.019		
un	0.012		
faites	0.011		

TABLE 9

Translation and fertility probabilities for should. should			
f	t(f e)	φ	n(φ e)
devrait	0.330	1	0.649
devraient	0.123	0	0.336
devrions	0.109	2	0.014
faudrait	0.073		
faud	0.058		
doit	0.058		
aurait	0.041		
doivent	0.024		
devons	0.017		
devrais	0.013		

TABLE 10

Translation and fertility probabilities for national. national			
f	t(f e)	φ	n(φ e)
nationale	0.469	1	0.905
national	0.418	0	0.094
nationaux	0.054		
nationales	0.029		

TABLE 11

Translation and fertility probabilities for the. the			
f	t(f e)	φ	n(φ e)
le	0.497	1	0.746
la	0.207	0	0.254
les	0.155		
l'	0.086		
ce	0.018		
cette	0.011		

TABLE 12

Translation and fertility probabilities for farmers. farmers			
f	t(f e)	φ	n(φ e)
agriculteurs	0.442	2	0.731
les	0.418	1	0.228
cultivateurs	0.046	0	0.039
producteurs	0.021		

63

TABLE 13

Translation and fertility probabilities for external. external			
f	t(f e)	$\phi$	n( $\phi$  e)
extérieures	0.944	1	0.967
extérieur	0.015	0	0.028
externe	0.011		
extérieurs	0.010		

TABLE 14

Translation and fertility probabilities for answer. answer			
f	t(f e)	$\phi$	n( $\phi$  e)
réponse	0.442	1	0.809
répondre	0.223	2	0.115
répondu	0.041	0	0.074
a	0.038		
solution	0.027		
répondez	0.021		
répondrai	0.016		
reponde	0.014		
y	0.013		
ma	0.010		

TABLE 15

Translation and fertility probabilities for oil. oil			
f	t(f e)	$\phi$	n( $\phi$  e)
pétrole	0.558	1	0.760
petrolières	0.138	0	0.181
pétrolière	0.109	2	0.057
le	0.054		
pétrolier	0.030		
pétroliers	0.024		
huile	0.020		
Oil	0.013		

TABLE 16

Translation and fertility probabilities for former. former			
f	t(f e)	$\phi$	n( $\phi$  e)
ancien	0.592	1	0.866
anciens	0.092	0	0.074
ex	0.092	2	0.060
précédent	0.054		
l'	0.043		
ancienne	0.018		
été	0.013		

TABLE 17

Translation and fertility probabilities for not. not			
f	t(f e)	$\phi$	n( $\phi$  e)
ne	0.497	2	0.735
pas	0.442	0	0.154
non	0.029	1	0.107
rien	0.011		

together with an article. Thus, farmers typically has a

64

fertility 2 and usually produces either agriculteurs or les. Additional examples are included in Tables 13 through 17, which show the translation and fertility probabilities for external, answer, oil, farmer, and not.

FIGS. 37, 38, and 39 show automatically derived alignments for three translations. In the terminology used above, each alignment is  $\delta^-(V(f|e); 2)$ . It should be understood that these alignments have been found by a process that involves no explicit knowledge of either French or English. Every fact adduced to support them has been discovered automatically from the 1,778,620 translations that constitute the training data. This data, in turn, is the product of a process the sole linguistic input of which is a set of rules explaining how to find sentence boundaries in the two languages. It may justifiably be argued, therefore, that these alignments are inherent in the Canadian Hansard data itself.

In the alignment shown in FIG. 37, all but one of the English words has fertility 1. The final prepositional phrase has been moved to the front of the French sentence, but otherwise the translation is almost verbatim. Notice, however, that the new proposal has been translated into les nouvelles propositions, demonstrating that number is not an invariant under translation. The empty notion has fertility 5 here. It generates  $en_1$ ,  $de_3$ , the comma,  $de_{16}$ , and  $de_{18}$ .

In FIG. 38, two of the English words have fertility 0, one has fertility 2, and one, embattled, has fertility 5. Embattled is another word, like nodding, that eludes the French grasp and comes with a panoply of multi-word translations.

The final example, in FIG. 39, has several features that bear comment. The second word, Speaker, is connected to the sequence l'Orateur. Like farmers above, it has trained to produce both the word that one naturally thinks of as its translation and the article in front of it. In the Hansard data, Speaker always has fertility 2 and produces equally often l'Orateur and le président. Later in the sentence, starred is connected to the phrase marquées de un astérisque. From an initial situation in which each French word is equally probable as a translation of starred, the system has arrived, through training, at a situation where it is able to connect to just the right string of four words. Near the end of the sentence, give is connected to donnerai, the first person singular future of donner, which means to give. It might be better if both will and give were connected to donnerai, but by limiting notions to no more than one word, Model 5 precludes this possibility. Finally, the last twelve words of the English sentence, I now have the answer and will give it to the House, clearly correspond to the last seven words of the French sentence, je donnerai la réponse à la Chambre, but, literally, the French is I will give the answer to the House. There is nothing about not, have, and, or it, and each of these words has fertility 0. Translations that are as far as this from the literal are rather more the rule than the exception in the training data.

## 9.12 Better Translation Models

Models 1 through 5 provide an effective means for obtaining word-by-word alignments of translations, but as a means to achieve the real goal of translation, there is room for improvement.

Two shortcomings are immediately evident. Firstly, because they ignore multi-word notions, these models are forced to give a false, or at least an unsatisfactory, account of some features in many translations. Also, these models are deficient, either in fact, as with Models 3 and 4, or in spirit, as with Models 1, 2, and 5.

## 9.12.1 Deficiency

It has been argued above that neither spiritual nor actual deficiency poses a serious problem, but this is not entirely true. Let  $w(e)$  be the sum of  $P(f|e)$  over well-formed French strings and let  $i(e)$  be the sum over ill-formed French strings. In a deficient model,  $w(e)+i(e)<1$ . In this case, the remainder of the probability is concentrated on the event failure and so  $w(e)+i(e)+P(\text{failure})=1$ . Clearly, a model is deficient precisely when  $P(\text{failure})>0$ . If  $P(\text{failure})=0$ , but  $i(e)>0$ , then the model is spiritually deficient. If  $w(e)$  were independent of  $e$ , neither form of deficiency would pose a problem, but because Models 1-5 have no long-term constraints,  $w(e)$  decreases exponentially with  $l$ . When computing alignments, even this creates no problem because  $e$  and  $f$  are known. However, for a given  $f$ , if the goal is to discover the most probable  $\hat{e}$ , then the product  $P(e)P(f|e)$  is too small for long English strings as compared with short ones. As a result, short English strings are improperly favored over longer English strings. This tendency is counteracted in part by the following modification:

- Replace  $P(f|e)$  with  $cP(f|e)$  for some empirically chosen constant  $c$ .

This modification is treatment of the symptom rather than treatment of the disease itself, but it offers some temporary relief. The cure lies in better modelling.

## 9.12.2 Multi-word Notions

Models 1 through 5 all assign non-zero probability only to alignments with notions containing no more than one word each. Except in Models 4 and 5, the concept of a notion plays little rôle in the development. Even in these models, notions are determined implicitly by the fertilities of the words in the alignment: words for which the fertility is greater than zero make up one-word notions; those for which it is zero do not. It is not hard to give a method for extending the generative process upon which Models 3, 4, and 5 are based to encompass multi-word notions. This method comprises the following enhancements:

- Including a step for selecting the notional scheme;
- Ascribing fertilities to notions rather than to words;
- Requiring that the fertility of each notion be greater than zero.
- In Equation (71), replacing the products over words in an English string with products over notions in the notional scheme.

Extensions beyond one-word notions, however, requires some care. In an embodiment as described in the subsection Results, an English string can contain any of 42,005 one-word notions, but there are more than 1.7 billion possible two-word notions, more than 74 trillion three-word notions, and so on. Clearly, it is necessary to be discriminating in choosing potential multi-word notions. The caution displayed by Models 1-5 in limiting consideration to notions with fewer than two words is motivated primarily by respect for the featureless desert that multi-word notions offer a priori. The Viterbi alignments computed with Model 5 give a frame of reference from which to develop methods that expands horizons to multi-word notions. These methods involve:

- Promoting a multi-word sequence to notionhood when its translations as a unit, as observed in the Viterbi alignments, differ substantially from what is expected on the basis of its component parts.

In English, for example, either a boat or a person can be left high and dry, but in French, un bateau is not left haut et sec, nor une personne haute et sèche. Rather, a boat is left

échoué and a person en plan. High and dry, therefore, is a promising three-word notion because its translation is not compositional.

## 10 MATHEMATICAL SUMMARY OF TRANSLATION MODELS

## 10.1 Summary of Notation

$\mathcal{E}$	English vocabulary
$e$	English word
$c$	English sentence
$l$	length of $e$
$i$	position in $e$ , $i = 0, 1, \dots, l$
$e_i$	word $i$ of $e$
$e_0$	the empty notion
$c_i$	$e_1 e_2 \dots e_i$
$\mathcal{F}$	French vocabulary
$f$	French word
$m$	French sentence
$j$	length of $f$
$f_j$	position in $f$ , $j = 1, 2, \dots, m$
$f_j^i$	word $j$ of $f$
$a$	$f_1 f_2 \dots f_j$
$a_j$	alignment
$a_i^j$	position in $e$ connected to position $j$ of $f$ for alignment $a$
$\Phi_i$	$a_1 a_2 \dots a_i$
$\Phi_i^j$	number of positions of $f$ connected to position $i$ of $e$
$\tau$	$\Phi_1 \Phi_2 \dots \Phi_i$
$\tau_i$	tableau - a sequence of tablets, where a tablet is a sequence of French word
$\tau_0 \tau_1 \dots \tau_i$	tablet $i$ of $\tau$
$\Phi_i$	$\tau_0 \tau_1 \dots \tau_i$
$k$	length of $\tau_i$
$\tau_{ik}$	position within a tablet, $k = 1, 2, \dots, \Phi_i$
$\pi$	word $k$ of $\tau_i$
$\pi_{ik}$	a permutation of the positions of a tableau
$\pi_{i1} \pi_{i2} \dots \pi_{ik}$	position in $f$ for word $k$ of $\tau_i$ for permutation $\pi$
$V(f e)$	$\pi_{i1} \pi_{i2} \dots \pi_{ik}$
$V_{i \leftarrow j}(f e)$	Viterbi alignment for $(f e)$
$\mathcal{N}(a)$	Viterbi alignment for $(f e)$ with $ij$ pegged
$\mathcal{N}_{ij}(a)$	neighboring alignments of $a$
$b(a)$	neighboring alignments of $a$ with $ij$ pegged
$b^-(a)$	alignment in $\mathcal{N}(a)$ with greatest probability
$b_{i \leftarrow j}(a)$	alignment obtained by applying $b$ repeatedly to $a$
$b_{i \leftarrow j}^-(a)$	alignment in $\mathcal{N}_{ij}(a)$ with greatest probability
$A(e)$	alignment obtained by applying $b_{i \leftarrow j}$ repeatedly to $a$
$\mathcal{B}(f)$	class of English word $e$
$\Delta_j$	class of French word $f$
$v, v'$	displacement of a word in $f$
$p_i$	vacancies in $f$
$c_i$	first position in $e$ to the left of $i$ that has non-zero fertility
$\{i\}$	average position in $f$ of the words connected to position $i$ of $e$
$O_i$	position in $e$ of the $i^{\text{th}}$ one word notion
$P_\theta$	$c_{\{i\}}$
$C(f, e)$	translation model $P$ with parameter values $\theta$
$\psi(P_\theta)$	empirical distribution of a sample
$R(P_\theta, P_\theta)$	log-likelihood objective function
$t(f e)$	relative objective function
$\epsilon(\text{null})$	translation probabilities (All Models)
$n(\phi_k)$	sentence length probabilities (Models 1 and 2)
$p_0, p_1$	fertility probabilities (Models 3, 4, and 5)
	fertility probabilities for $e_0$ (Models 3, 4,

-continued

$a(i j, l, m)$	and 5)
$d(j i, l, m)$	alignment probabilities (Model 2)
$d_1(\Delta j A, B)$	distortion probabilities (Model 3)
	distortion probabilities for the first word
$d_{>1}(\Delta j B)$	of a tablet (Model 4)
	distortion probabilities for the other words
$d_{>1}(l\_or\_r B, v, v')$	of a tablet (Model 4)
	distortion probabilities for choosing left or
$d_L(\Delta j B, v)$	right movement (Model 5)
	distortion probabilities for leftward
$d_R(\Delta j B, v)$	movement of the first word of a tablet
	(Model 5)
$d_{>1}(\Delta j B, v)$	distortion probabilities for rightward
	movement of the first word of a tablet
	(Model 5)
	distortion probabilities for movement
	of the other words of a tablet (Model 5)

## 10.2 Model 1

## Parameters

 $\epsilon(ml)$  string length probabilities $t(f|e)$  translation probabilitiesHere  $f \in \mathcal{F}$ ;  $e \in \mathcal{E}$  or  $e = e_0$ ;  $l = 1, 2, \dots$ ; and  $m = 1, 2, \dots$ 

## General Formula

$$P_\theta(f, a|e) = P_\theta(m|e) P_\theta(a|m, e) P_\theta(f|a, m, e) \quad (92)$$

## Assumptions

$$P_\theta(m|e) = \epsilon(ml) \quad (93)$$

$$P_\theta(a|m, e) = (l+1)^{-m} \quad (94)$$

$$P_\theta(f|a, m, e) = \prod_{j=1}^m t(f_j|e_{a_j}) \quad (95)$$

This model is not deficient.

## Generation

Equations (92)–(95) describe the following process for producing  $f$  from  $e$ :

1. Choose a length  $m$  for  $f$  according to the probability distribution  $\epsilon(ml)$ .
2. For each  $j = 1, 2, \dots, m$ , choose  $a_j$  from  $0, 1, 2, \dots, l$  according to the uniform distribution.
3. For each  $j = 1, 2, \dots, m$ , choose a French word  $f_j$  according to the distribution  $t(f_j|e_{a_j})$ .

## Useful Formulae

Because of the independence assumptions (93)–(95), the sum over alignments (26) can be calculated in closed form:

$$P_\theta(f|e) = \sum_a P_\theta(f, a|e) \quad (96)$$

$$= \epsilon(ml) (l+1)^{-1} \sum_{a_1=0}^l \dots \sum_{a_m=0}^l \prod_{j=1}^m t(f_j|e_{a_j}) \quad (97)$$

-continued

$$= \epsilon(ml) (l+1)^{-m} \pi \sum_{j=1}^m \prod_{i=0}^l t(f_j|e_i). \quad (98)$$

Equation (98) is useful in computations since it involves only  $O(lm)$  arithmetic operations, whereas the original sum over alignments (97) involves  $O(l^m)$  operations.

## Concavity

The objective function (25) for this model is a strictly concave function of the parameters. In fact, from Equations (25) and (98),

$$\psi(P_\theta) = \quad (99)$$

$$\sum_{f,e} C(f,e) \sum_{j=1}^m \log \sum_{i=0}^l t(f_j|e_i) + \sum_{f,e} C(f,e) \log \epsilon(ml) + \text{constant}$$

which is clearly concave in  $\epsilon(ml)$  and  $t(f|e)$  since the logarithm of a sum is concave, and the sum of concave function is concave.Because  $\psi$  is concave, it has a unique local maximum. Moreover, this maximum can be found using the method of FIG. 28, provided that none of the initial parameter values is zero.

## 10.3 Model 2

## Parameters

 $\epsilon(ml)$  string length probabilities $t(f|e)$  translation probabilities $a(i|j, l, m)$  alignment probabilitiesHere  $i = 0, \dots, l$ ; and  $j = 1, \dots, m$ .

## General Formula

$$P_\theta(f, a|e) = P_\theta(m|e) P_\theta(a|m, e) P_\theta(f|a, m, e) \quad (100)$$

## Assumptions

$$P_\theta(m|e) = \epsilon(ml) \quad (101)$$

$$P_\theta(a|m, e) = \prod_{j=1}^m a(a_j|j, l, m) \quad (102)$$

$$P_\theta(f|a, m, e) = \prod_{j=1}^m t(f_j|e_{a_j}) \quad (103)$$

This model is not deficient. Model 1 is the special case of this model in which the alignment probabilities are uniform:  $a(i|j, l, m) = (l+1)^{-1}$  for all  $i$ .

## Generation

Equations (100)–(103) describe the following process for producing  $f$  from  $e$ :

1. Choose a length  $m$  for  $f$  according to the distribution  $\epsilon(ml)$ .
2. For each  $j = 1, 2, \dots, m$ , choose  $a_j$  from  $0, 1, 2, \dots, l$  according to the distribution  $a(a_j|j, l, m)$ .
3. For each  $j$ , choose a French word  $f_j$  according to the distribution  $t(f_j|e_{a_j})$ .



## Useful Formulae

Just as for Model 1, the independence assumptions allow us to calculate the sum over alignments (26) in closed form:

$$P_{\theta}(fle) = \sum_a P_{\theta}(f, a|e) \quad (104)$$

$$= \epsilon(m|l) (l+1)^{-1} \sum_{a_1=0}^l \dots \sum_{a_m=0}^m \pi \tau(f|e_{a_1}) \alpha(a|f, l, m) \quad (105)$$

$$= \epsilon(m|l) (l+1)^{-m} \pi \sum_{j=1}^m \tau(f|e_j) \alpha(i|j, l, m). \quad (106)$$

By assumption (102) the connections of  $a$  are independent given the length  $m$  of  $f$ . From Equation (106) it follows that they are also independent given  $f$ :

$$P_{\theta}(a|f, e) = \pi \prod_{j=1}^m p_{\theta}(a|j, f, e). \quad (107)$$

where

$$p_{\theta}(i|j, f, e) = \frac{\gamma(i, j, f, e)}{\sum_{i'} \gamma(i', j, f, e)} \quad (108)$$

with

$$\gamma(i, j, f, e) = \tau(f|e_i) \alpha(i|j, l, m). \quad (109)$$

## Viterbi Alignment

For this model, and thus also for Model 1, the Viterbi alignment  $V(f|e)$  between a pair of strings  $(f, e)$  can be expressed in closed form:

$$V(f|e) = \arg \max_i (f|e_i) \alpha(i|j, l, m). \quad (109)$$

## Parameter Reestimation Formulae

The parameter values  $\theta$  that maximize the relative objective function  $R(\hat{P}_{\theta}, P_{\theta})$  can be found by applying the considerations of subsection 8.3.4. The counts  $c(\omega; a, f, e)$  of Equation (38) are

$$c(f|e; a, f, e) = \sum_{j=1}^m \delta(e, e_j) \delta(f, f_j), \quad (110)$$

$$c(i|j, l, m; a, f, e) = \delta(i, a_j). \quad (111)$$

The parameter reestimation formulae for  $t(f|e)$  and  $a(i|j, l, m)$  are obtained by using these counts in Equations (42)–(46).

Equation (44) requires a sum over alignments. If  $\hat{P}_{\theta}$  satisfies

$$P_{\theta}(a|f, e) = \pi \prod_{j=1}^m \hat{p}_{\theta}(a|j, f, e), \quad (112)$$

as is the case for Models 1 and 2 (see Equation (107)), then this sum can be calculated explicitly:

$$\hat{c}_{\theta}(f|e; a, f, e) = \sum_{i=0}^l \sum_{j=1}^m \hat{p}_{\theta}(i|j, f, e) \delta(e, e_i) \delta(f, f_j), \quad (113)$$

$$\hat{c}_{\theta}(i|j, f, e) = \hat{p}_{\theta}(i|j, f, e). \quad (114)$$

Equations (110)–(114) involve only  $O(lm)$  arithmetic operations, whereas the sum over alignments involves  $O(l^m)$  operations.

## 10.4 Model 3

## Parameters

$t(f e)$	translation probabilities
$n(\phi e)$	fertility probabilities
$p_0, p_1$	fertility probabilities for $e_0$
$d(i l, m)$	distortion probabilities

Here  $\phi = 0, 1, 2, \dots$

## General Formulae

$$P_{\theta}(\tau, \pi|e) = P_{\theta}(\phi|e) P_{\theta}(\tau|\phi, e) P_{\theta}(\pi|\tau, \phi, e) \quad (115)$$

$$P_{\theta}(f, a|e) = \sum_{(\tau, \pi) \in \langle f, a \rangle} P_{\theta}(\tau, \pi|e) \quad (116)$$

Here  $\langle f, a \rangle$  is the set of all  $(\tau, \pi)$  consistent with  $(f, a)$ :

$$(\tau, \pi) \in \langle f, a \rangle \text{ if for all } i=0, 1, \dots, l \text{ and } k=1, 2, \dots, \phi_i, f_{i+k} = \tau_{i+k} \text{ and } a_{i+k} = i. \quad (117)$$

## Assumptions

$$P_{\theta}(\phi|e) = n_0 \left( \phi_0! \sum_{i=1}^l \phi_i \right) \frac{1}{\pi} n(\phi|e_i) \quad (118)$$

$$P_{\theta}(\tau|\phi, e) = \frac{1}{\pi} \prod_{i=0}^l \prod_{k=1}^{\phi_i} t(\tau_{i+k}|e_i) \quad (119)$$

$$P_{\theta}(\pi|\tau, \phi, e) = \frac{1}{\phi_0!} \prod_{i=1}^l \prod_{k=1}^{\phi_i} \pi d(\pi_{i+k}|i, l, m) \quad (120)$$

where

$$n_0(\phi_0|m) = \left( \frac{m!}{\phi_0} \right) p_0^{m-\phi_0} p_1^{\phi_0}. \quad (121)$$

In Equation (120) the factor of  $1/\phi_0!$  accounts for the choices of  $\pi_{0k}$ ,  $k=1, 2, \dots, \phi_0$ . This model is deficient, since

$$P_{\theta}(\text{failure}|\tau, \phi, e) = 1 - \sum_{\pi} P_{\theta}(\pi|\tau, \phi, e) > 0. \quad (122)$$

## Generation

Equations (115)–(120) describe the following process for producing  $f$  or failure from  $e$ :

1. For each  $i=1, 2, \dots, l$ , choose a length  $\phi_i$  for  $\tau_i$  according to the distribution  $n(\phi_i|e_i)$ .
2. Choose a length  $\phi_0$  for  $\tau_0$  according to the distribution  $n_0(\phi_0|\sum_{i=1}^l \phi_i)$ .
3. Let  $m = \phi_0 + \sum_{i=1}^l \phi_i$ .
4. For each  $i=0, 1, \dots, l$  and each  $k=1, 2, \dots, \phi_i$ , choose a French word  $\tau_{i+k}$  according to the distribution  $t(\tau_{i+k}|e_i)$ .
5. For each  $i=1, 2, \dots, l$  and  $k=1, 2, \dots, \phi_i$ , choose a position  $\pi_{i+k}$  from  $1, \dots, m$  according to the distribution  $d(\pi_{i+k}|i, l, m)$ .

6. If any position has been chosen more than once then return failure.

7. For each  $k=1, 2, \dots, \phi_0$ , choose a position  $\pi_{\alpha k}$  from the  $\phi_0-k+1$  remaining vacant positions in  $1, 2, \dots, m$  according to the uniform distribution.

8. Let  $f$  be the string with  $f_{\pi k} = \tau_{ik}$ .

#### Useful Formulae

From Equations (118)–(120) it follows that if  $(\tau, \pi)$  are consistent with  $(f, a)$  then

$$P_0(\tau|f, e) = \sum_{j=1}^m \tau(f|e_{\alpha_j}), \quad (123)$$

$$P_0(\pi|\tau, \phi, e) = \frac{1}{\phi_0!} \prod_{j: a_j \neq 0} \pi d(j|a_j, l, m). \quad (124)$$

In Equation (124), the product runs over all  $j=1, 2, \dots, m$  except those for which  $a_j=0$ .

By summing over all pairs  $(\tau, \pi)$  consistent with  $(f, a)$  it follows that

$$\begin{aligned} P_0(f, a|e) &= \sum_{(\tau, \pi) \in \langle f, a \rangle} P_0(\tau, \pi|e) \\ &= n_0 \left( \phi_0! \sum_{i=1}^l \phi_i \right) \prod_{i=1}^l \pi n(\phi_i|e_i) \prod_{j=1}^m \tau(f|e_{\alpha_j}) \prod_{j: a_j \neq 0} \pi d(j|a_j, l, m). \end{aligned} \quad (125)$$

The factors of  $\phi_i!$  in Equation (126) arise because there are  $\prod_{i=0}^l \phi_i!$  equally probable terms in the sum (125).

#### Parameter Reestimation Formulae

The parameter values  $\theta$  that maximize the relative objective function  $R(\tilde{P}_\theta, P_\theta)$  can be found by applying the considerations of Section 8.3.4. The counts  $c(\omega; a, f, e)$  of Equation (38) are

$$c(f|e; a, f, e) = \sum_{j=1}^m \delta(e, e_{\alpha_j}) \delta(f, f_j), \quad (127)$$

$$c(j|i, l, m; a, f, e) = \delta(i, a_j), \quad (128)$$

$$c(\phi|e; a, f, e) = \sum_{i=1}^l \delta(e, e_i) \delta(\phi, \phi_i). \quad (129)$$

The parameter reestimation formulae for  $t(f|e)$ ,  $a(j|i, l, m)$ , and  $t(\phi|e)$  are obtained by using these counts in Equations (42)–(46).

Equation (44) requires a sum over alignments. If  $(\tilde{P}_\theta)$  satisfies

$$P_\theta(a|f, e) = \prod_{j=1}^m \tilde{P}_\theta(a_j|f, e), \quad (130)$$

is the case for Models 1 and 2 (see Equation (107)), then this sum can be calculated explicitly for  $\tilde{c}_\theta(f|e, f, e)$  and  $\tilde{c}_\theta(j|i, f, e)$ :

$$\tilde{c}_\theta(f|e; f, e) = \sum_{i=0}^l \sum_{j=1}^m \tilde{P}_\theta(i|j, f, e) \delta(e, e_i) \delta(f, f_j), \quad (131)$$

-continued

$$\tilde{c}_\theta(j|i, f, e) = \tilde{P}_\theta(j|i, f, e). \quad (132)$$

Unfortunately, there is no analogous formula for  $\tilde{c}_\theta(\phi|e; f, e)$ . Instead the following formulae hold:

$$\tilde{c}_\theta(\phi|e; f, e) = \sum_{i=1}^l \delta(e, e_i) \prod_{j=1}^m (1 - \tilde{P}_\theta(j|i, f, e)) \quad (133)$$

$$\tilde{P}_\theta(j|i, f, e) = \sum_{\gamma \in \Gamma_\phi} \frac{\alpha_\gamma(f, e)^\gamma}{\gamma!}.$$

$$\alpha_\gamma(f, e) = \frac{(-1)^{k+1}}{k!} \frac{1}{k} \sum_{j=1}^m \tilde{P}_\theta(j|f, e)^k, \quad (134)$$

$$\tilde{P}_\theta(j|f, e) = \frac{\tilde{P}_\theta(j|f, e)}{1 - \tilde{P}_\theta(j|f, e)}. \quad (135)$$

In Equation (133),  $\Gamma_\phi$  denotes the set of all partitions of  $\phi$ .

$$(125)$$

$$(126)$$

Recall that a partition of  $\phi$  is a decomposition of  $\phi$  as a sum of positive integers. For example,  $\phi=5$  has 7 partitions since  $1+1+1+1+1=1+1+1+2=1+1+3=1+2+2=1+4=2+3=5$ .

For a partition  $\gamma$ , let  $\gamma_k$  be the number of times that  $k$  appears in the sum, so that  $\phi = \sum_{k=1}^{\phi} \gamma_k k$ . If  $\gamma$  is the partition corresponding to  $1+1+3$ , then  $\gamma_1=2$ ,  $\gamma_3=1$ , and  $\gamma_k=0$  for  $k$  other than 1 or 3. Here the convention that  $\Gamma_0$  consists of the single element  $\gamma$  with  $\gamma_k=0$  for all  $k$  has been adopted.

Equation (133) allows us to compute the counts  $\tilde{c}_\theta(\phi|e; f, e)$  in  $O(lm+\phi g)$  operations, where  $g$  is the number of partitions of  $\phi$ . Although  $g$  grows with  $\phi$  like  $(4\sqrt{3}\phi)^{-1} \exp \pi\sqrt{2\phi/3}$ , it is manageably small for small  $\phi$ . For example,  $\phi=10$  has 42 partitions.

#### Proof of Formula (133)

Introduce the generating functional

$$G(x|e, f, e) = \sum_{\phi=0}^{\infty} \tilde{c}_\theta(\phi|e; f, e) x^\phi, \quad (136)$$

where  $x$  is an indeterminant. Then

$$G(x|e, f, e) = \sum_{\phi=0}^{\infty} \sum_{\alpha_1=0}^l \dots \quad (137)$$

$$\sum_{\alpha_m=0}^m \prod_{j=1}^m \tilde{P}_\theta(\alpha_j|j, f, e) \sum_{i=1}^l \delta(e, e_i) \delta(\phi, \phi_i) x^\phi$$

$$= \sum_{i=1}^l \delta(e, e_i) \sum_{\alpha_1=0}^l \dots \sum_{\alpha_m=0}^m \prod_{j=1}^m \tilde{P}_\theta(\alpha_j|j, f, e) x^{\sum \alpha_j} \quad (138)$$

-continued

$$= \sum_{i=1}^l \delta(e, e_i) \sum_{\alpha_1=0}^l \dots \sum_{\alpha_m=0}^l \pi \bar{p}_\theta(\alpha_j | f, e) x^{k(i, \alpha_j)} \quad (139)$$

$$= \sum_{i=1}^l \delta(e, e_i) \pi \sum_{j=1}^m \sum_{\alpha=0}^l \bar{p}_\theta(\alpha | j, f, e) x^{k(i, \alpha)} \quad (140)$$

$$= \sum_{i=1}^l \delta(e, e_i) \pi \sum_{j=1}^m (1 - \bar{p}_\theta(i | j, f, e) + x \bar{p}_\theta(i | j, f, e)) \quad (141)$$

$$= \sum_{i=1}^l \delta(e, e_i) \pi \sum_{j=1}^m (1 - \bar{p}_\theta(i | j, f, e)) \pi \sum_{j=1}^m (1 + \beta_j(f, e) x) \quad (142)$$

To obtain Equation (138) rearrange the order of summation and sum over  $\phi$  to eliminate the  $\delta$ -function of  $\phi$ . To obtain Equation (139), note that  $\phi = \sum_{j=1}^m \delta(i, a_j)$  and so  $x^{\phi} = \prod_{j=1}^m x^{\delta(i, a_j)}$ . To obtain Equation (140), interchange the order of the sums on  $a_j$  and the product on  $j$ . To obtain Equation (141), note that in the sum on  $a$ , the only term for which the power of  $x$  is nonzero is the one for which  $a=i$ .

Now note that for any indeterminates  $x, y_1, y_2, \dots, y_m$ ,

$$\pi \sum_{j=1}^m (1 + y_j x) = \sum_{\phi=0}^m x^{\phi} \sum_{\gamma \in T_{\phi}} \pi \frac{z_k^{\gamma_k}}{\gamma_k!} \quad (143)$$

$$\text{where } z_k = \frac{(-1)^{k+1}}{k} \sum_{j=1}^m (y_j)^k. \quad (144)$$

This follows from the calculation

$$\pi \sum_{j=1}^m (1 + y_j x) = \exp \sum_{j=1}^m \log(1 + y_j x) = \quad (145)$$

$$\exp \sum_{j=1}^m \sum_{k=1}^{\infty} \frac{(-1)^{k+1} (y_j x)^k}{k} = \exp \sum_{k=1}^{\infty} z_k x^k = \sum_{n=0}^{\infty} \frac{1}{n!} \left( \sum_{k=1}^{\infty} z_k x^k \right)^n \quad (146)$$

$$= \sum_{n=0}^{\infty} \frac{1}{n!} \sum_{\gamma_1 \gamma_2 \dots} \left( \begin{matrix} n \\ \gamma_1 \gamma_2 \dots \end{matrix} \right) \pi \sum_{k=1}^{\infty} (z_k x^k)^{\gamma_k} = \quad (147)$$

$$\sum_{\phi=0}^{\infty} x^{\phi} \sum_{\gamma \in T_{\phi}} \pi \frac{z_k^{\gamma_k}}{\gamma_k!} \quad (148)$$

The reader will notice that the left-hand side of Equation (145) involves only powers of  $x$  up to  $m$ , while the subsequent Equations (146)–(147) involve all powers of  $x$ . This is because the  $z_k$  are not algebraically independent. In fact, for  $\phi > m$ , the coefficient of  $x^{\phi}$  on the right-hand side of Equation (147) must be zero. It follows that  $z_k$  can be expressed as a polynomial in  $z_1, z_2, \dots, z_m$ .

Equation (143) can be used to identify the coefficient of  $x^{\phi}$  in Equation (142). Equation (133) followed by combining Equations (142), (143), and the Definitions (134)–(136) and (144).

## 10.5 Model 4

## Parameters

5	$t(f e)$	translation probabilities
	$n(\phi e)$	fertility probabilities
	$p_0, p_1$	fertility probabilities for $e_0$
10	$d_1(\Delta_j   \mathcal{A}, \mathcal{B})$	distortion probabilities for movement of the first word of a tablet
	$d_{<1}(\Delta_j   \mathcal{B})$	distortion probabilities for movement of other words of a tablet

Here  $\Delta_j$  is an integer;  $\mathcal{A}$  is an English class; and  $\mathcal{B}$  is a French class.

## General Formulae

$$P_\theta(\tau, \pi | e) = P_\theta(\phi | e) P_\theta(\tau | \phi, e) P_\theta(\pi | \tau, \phi, e) \quad (148)$$

$$P_\theta(f, a | e) = \sum_{(\tau, \pi) \in \mathcal{C}(f, a)} P_\theta(\tau, \pi | e) \quad (149)$$

## Assumptions

$$P_\theta(\phi | e) = n_0 \left( \phi_0! \prod_{i=1}^l \phi_i! \right) \pi \sum_{i=1}^l n(\phi_i | e_i) \quad (150)$$

$$P_\theta(\tau | \phi, e) = \pi \prod_{i=0}^l \pi \tau(\tau_{ik} | e_i) \quad (151)$$

$$P_\theta(\pi | \tau, \phi, e) = \frac{1}{\phi_0!} \prod_{i=1}^l \pi \pi p_{ik}(\pi_{ik}) \quad (152)$$

where

$$n_0(\phi_0 | m) = \left( \begin{matrix} m \\ \phi_0 \end{matrix} \right) p_0^{m-\phi_0} p_1^{\phi_0} \quad (153)$$

$$p_{ik}(j) = \begin{cases} d_1(j - c_{p_i} | \mathcal{A}(e_{p_i}), \mathcal{B}(\tau_{i1})) & \text{if } k=1 \\ d_{>1}(j - \tau_{ik-1} | \mathcal{B}(\tau_{ik})) & \text{if } k > 1 \end{cases} \quad (154)$$

In Equation (154),  $p_i$  is the first position to the left of  $i$  for which  $\phi_i > 0$ , and  $c_p$  is the ceiling of the average position of the words of  $\tau_p$ .

$$p_i = \max_{i' < i} \{i' : \phi_{i'} > 0\}, c_p = \lceil \phi_p^{-1} \sum_{k=1}^{\phi_p} \tau_{pk} \rceil \quad (155)$$

This model is deficient, since

$$P_\theta(\text{failure} | \tau, \phi, e) = 1 - \sum_{\pi} P_\theta(\pi | \tau, \phi, e) > 0. \quad (156)$$

Note that Equations (150), (151), and (153) are identical to the corresponding formulate (118), (119), and (121) for Model 3.

## Generation

Equations (148)–(152) describe the following process for producing  $f$  or failure from  $e$ : 1.–4. Choose a tableau  $\tau$  by following Steps 1–4 for Model 3.

5. For each  $i=1, 2, \dots, l$  and each  $k=1, 2, \dots, \phi_i$  choose a position  $\pi_{ik}$  as follows.

If  $k=1$  then choose  $\pi_{i1}$  according to the distribution

$$d_1(\pi_{i1}-c_{pi})! \mathcal{B}(\tau_{i1}). \quad 5$$

If  $k>1$  then choose  $\pi_{ik}$  greater than  $\pi_{ik-1}$  according to the distribution  $d_{>1}(\pi_{ik}-\pi_{ik-1})! \mathcal{B}(\pi_{ik}).$

6-8. Finish generating  $f$  by following Steps 6-8 for Model 3. 10

10.6 Model 5

### Parameters

$r(f e)$	translation probabilities
$n(\phi e)$	fertility probabilities
$p_{0,p_1}$	fertility probabilities for $e_0$
$d_L(\Delta j) \mathcal{B}_{,v}$	distortion probabilities for leftward movement of the first word of a tablet
$d_R(\Delta j) \mathcal{B}_{,v}$	distortion probabilities for rightward movement of the first word of a tablet
$d_{LR}(l\_or\_r) \mathcal{B}_{,v,v'}$	distortion probabilities for choosing left or right movement
$d_{>1}(\Delta j) \mathcal{B}_{,v}$	distortion probabilities for movement of other words of a tablet

Here  $v, v'=1, 2, \dots, m$ ; and  $l\_or\_r$ =left or right.

### General Formula

$$P_\theta(\tau, f|e) = P_\theta(\phi|e) P_\theta(\tau|\phi, e) P_\theta(\pi|\tau, \phi, e) \quad (157)$$

$$P_\theta(f, a|e) = \sum_{(\tau, \pi) \in \mathcal{C}_{f,a}} P_\theta(\tau, f|e) \quad (159)$$

### Assumptions

$$P_\theta(\phi|e) = n_0 \left( \phi_0! \sum_{i=1}^l \phi_i \right) \prod_{i=1}^l \pi_i n(\phi_i|e) \quad (159)$$

$$P_\theta(\tau|\phi, e) = \prod_{i=0}^l \prod_{k=1}^{\phi_i} r(\tau_{ik}|e_i) \quad (160)$$

$$P_\theta(\pi|\tau, \phi, e) = \frac{1}{\phi_0!} \prod_{i=1}^l \prod_{k=1}^{\phi_i} p_{ik}(\pi_{ik}) \quad (161)$$

where

$$n_0(\phi_0|m') = \left( \frac{m'}{\phi_0} \right) p_0^{m'-\phi_0} p_1^{\phi_0}.$$

-continued

(164)

$$p_{ik}(j) = \epsilon_{ik}(j) \left\{ \begin{array}{ll} d_{LR}(\text{right}) \mathcal{B}(\tau_{i1}, v_{i1}(c_{pi}), v_{i1}(m)) & \text{if } k=1 \text{ and } j \geq c_{pi} \\ d_R(v_{ik}(j) - v_{i1}(c_{pi})) \mathcal{B}(\tau_{i1}, v_{i1}(m) - v_{i1}(c_{pi})) & \text{if } k=1 \text{ and } j < c_{pi} \\ d_{LR}(\text{left}) \mathcal{B}(\tau_{i1}, v_{i1}(c_{pi}), v_{i1}(m)) & \text{if } k>1 \text{ and } j \geq c_{pi} \\ d_L(v_{i1}(c_{pi}) - v_{i1}(j)) \mathcal{B}(\tau_{i1}, v_{i1}(c_{pi})) & \text{if } k>1 \text{ and } j < c_{pi} \\ d_{>1}(v_{ik}(j) - v_{ik}(\pi_{ik-1})) \mathcal{B}(\tau_{ik}, v_{ik}(m) - v_{ik}(\pi_{ik-1})) & \text{if } k>1 \end{array} \right.$$

In Equation (164),  $p_i$  is the first position to the left of  $i$  which has a non-zero fertility; and  $c_p$  is the ceiling of the average position of the words of tablet  $p$  (see Equation (155)). Also,  $\epsilon_{ik}(j)$  is 1 if position  $j$  is vacant after all the words of tablets  $i'<i$  and the first  $k-1$  words of tablet  $i$  have been placed, and 0 otherwise.  $v_{ik}(j)$  is the number of vacancies not to the right of  $j$  at this time  $v_{ik}(j) = \sum_{j' \leq j} \epsilon_{ik}(j')$ .

This model is not deficient. Note that Equations (159), (160), and (163) are identical to the corresponding formulae for Model 3.

N.B. In the previous section, simplified embodiment of this model in which the probabilities  $d_{LR}$  do not appear is described.

### Generation

Equations (157)-(161) describe the following process for producing  $f$  from  $e$ :

1-4. Choose a tableau  $\tau$  by following Steps 1-4 for Model 3.

5. For each  $i=1, 2, \dots, l$  and each  $k=1, 2, \dots, \phi_i$  choose a position  $\pi_{ik}$  as follows:

If  $k=1$  then

(a) Choose left or right according to the distribution

$$d_{LR}(l\_or\_r) \mathcal{B}(\tau_{i1}, v_{i1}(c_{pi}), v_{i1}(m)).$$

(b) If right, then choose a vacant position  $\pi_{i1}$  greater than or equal to  $c_{pi}$  according to the distribution

$$d_R(v_{i1}(\pi_{i1}) - v_{i1}(c_{pi})) \mathcal{B}(\tau_{i1}, v_{i1}(m) - v_{i1}(c_{pi})).$$

(c) Otherwise, choose a vacant position  $\pi_{i1}$  less than or equal to  $c_{pi}$  according to the distribution

$$d_L(v_{i1}(c_{pi}) - v_{i1}(\pi_{i1})) \mathcal{B}(\tau_{i1}, v_{i1}(c_{pi})).$$

If  $k>1$  then choose a vacant position  $\pi_{ik}$  greater than  $\pi_{ik-1}$  according to the distribution  $d_{>1}(v_{ik}(\pi_{ik}) - v_{ik}(\pi_{ik-1}))$

$$\mathcal{B}(\tau_{ik}, v_{ik}(m) - v_{ik}(\pi_{ik-1})).$$

6-8. Finish generating  $f$  by following Steps 6-8 for Model 3.

### 11. SENSE DISAMBIGUATION

#### 11.1 Introduction

An alluring aspect of the statistical approach to machine translation is the systematic framework it provides for attacking the problem of lexical disambiguation. For example, an embodiment of the machine translation system depicted in FIG. 4 translates the French sentence Je vais prendre la décision as I will make the decision, correctly interpreting prendre as make. Its statistical translation

model, which supplies English translations of French words, prefers the more common translation take, but its trigram language model recognizes that the three-word sequence make the decision is much more probable than take the decision.

This system is not always so successful. It incorrectly renders *Je vais prendre ma propre décision* as *I will take my own decision*. Its language model does not realize that take my own decision is improbable because take and decision no longer fall within a single trigram.

Errors such as this are common because the statistical models of this system only capture local phenomena; if the context necessary to determine a translation falls outside the scope of these models, a word is likely to be translated incorrectly. However, if the relevant context is encoded locally, a word can be translated correctly.

As has been noted in Section 3, such encoding can be performed in the source-transduction phase 701 by a sense-labeling transducer.

In this section, the operation and construction of such a sense-labeling transducer is described. The goal of this transducer is to perform cross-lingual word-sense labeling. That is, this transducer labels words of a sentence in a source language so as to elucidate their translations into a target language. Such a transducer can also be used to label the words of an target sentence so as to elucidate their translations into a source language.

The design of this transducer is motivated by some examples. In some contexts the French verb *prendre* translates into English as *to take*, but in other contexts it translates as *to make*. A sense disambiguation transformation, by examining the contexts, might label occurrences of *prendre* that likely mean *to take* with one label, and other occurrences of *prendre* with another label. Then the uncertainty in the translation of *prendre* given the label would be less than the uncertainty in the translation of *prendre* without the label. Although the label does not provide any information that is not already present in the context, it encodes this information locally. Thus a local statistical model for the transfer of labeled sentences should be more accurate than one for the transfer of unlabeled ones.

While the translation of a word depends on many words in its context, it is often possible to obtain information by looking at only a single word. For example, in the sentence *Je vais prendre ma propre décision* (*I will make any own decision*), the verb *prendre* should be translated as *make* because its object is *décision* is replaced by *voiture* then *prendre* should be translated as *take*: *Je vais prendre ma propre voiture* (*I will take my own car*). Thus the uncertainty in the translation of *prendre* is reduced by asking a question about its object, which is often the first noun to its right. A sense can be assigned to *prendre* based upon the answer to this question.

As another example, in *Il doute que les nôtres gagnent* (*He doubts that we will win*), the word *il* should be translated as *he*. On the other hand, if *doute* is replaced by *faut* then *il* should be translated as *it*: *Il faut que les nôtres gagnent* (*It is necessary that we win*). Here, a sense label can be assigned to *il* by asking about the identity of the first verb to its right.

These examples motivate a sense-labeling scheme in which the label of a word is determined by a question about an informant word in its context. In the first example, the informant of *prendre* is the first noun to the right; in the second example, the informant of *il* is the first verb to the right. The first example is depicted in FIG. 40. The two sequence of this example are labeled 4001 and 4002 in the figure.

If more than two senses are desired for a word, then questions with more than two answers can be considered.

## 11.2 Design of a Sense-Labeling Transducer

FIG. 41 depicts an embodiment of a sense-labeling transducer based on this scheme. For expositional purposes, this embodiment will be discussed in the context of labeling words in a source language word sequence. It should be understood that in other embodiments, a sense-labeling transducer can accept as input more involved source-structure representations, including but not limited to, lexical morpheme sequences. In these embodiments, the constituents of a representation of a source word sequence are annotated by the sense-labeling transducer with labels that elucidate their translation into analogous constituents into a similar representation of a target word sequence. It should also be understood that in still other embodiments, a sense-labeling transducer annotates, target-structure representations, (not source-structure representations) with sense labels.

The operation of the embodiment of the sense-labeling transducer depicted in FIG. 41 comprises the steps of:

4101. Capturing an input sequence consisting of a sequence of words in a source language;

4102. For each word in the input sequence performing the Steps 4107, 4108, 4109, 4104 until no more words are available in Step 4105;

4107. For the input word being considered, finding a best informant site such as the noun to the right, the verb to the left, etc. A best informant for a word is obtained using a table 4103 stored in memory of informants and questions about the informants for every word in the source language vocabulary;

4108. Find the word at the best informant site in the input word sequence;

4109. Obtaining the class of the informant word as given by the answer to a question about the informant word;

4104. Labeling the original input word of the input sequence with the class of the informant word.

For the example depicted in FIG. 40, the informant site determined by Step 4107 is the noun to the right. For the first word sequence 4001 of this example, the informant word determined by Step 4108 is *décision*; for the second word sequence 4109, the informant word is *voiture*. In this example, the class of *décision* determined in Step 4109 is different than the class of *voiture*. Thus the label attached to *prendre* by Step 4104 is different for these two contexts of *prendre*.

## 11.3 Construction a Table of Informants and Questions

An important component of the sense-labeler depicted in FIG. 4104 is a table 4103 of informants and questions for every word in a source language vocabulary.

FIG. 42 depicts a method of constructing such a table. This method comprises the steps of:

4203. Performing the Steps 4201 and 4202 for each word in a source language vocabulary.

4201. For the word being considered, finding a good question for each of a plurality of informant sites. These informant sites are obtained from a table 4207 stored in memory of possible informant sites. Possible sites include but are not limited to, the nouns to the right and left, the verbs to the right and left, the words to the right and left, the words two positions to the right or left, etc. A method of finding a good question is described below. This method makes use of a table 4205 stored in memory probabilities derived from Viterbi alignments. These probabilities are also discussed below.

4202. Storing in a table 4206 of informants and questions, the informant site and the good question.

A method of carrying out the Step 4201 for finding a good question for each information site of a vocabulary word is depicted in FIG. 43. This method comprises the steps of:

4301. Performing the Steps 4302 and 4207 for each possible informant site. Possible informant sites are obtained from a table 4304 of such sites.

4302. For the informant site being considered, finding a question about informant words at this site that provides a lot of information about translations of the vocabulary word into the target language.

4207. Storing the vocabulary word, the informant site, and the good question in a table 4103.

#### 11.4 Mathematics of Constructing Questions

A method for carrying out the Step 4302 of finding a question about an informant is depicted in FIG. 44. In this subsection, the some preliminaries for describing this method are given. The notation used here is the same as that used in Sections 8–10.

##### 11.4.1 Statistical Translation with Transductions

Recall the setup depicted in FIG. 2. A system shown there employs

1. A source transducer 201 which encodes a source sentence  $f$  into an intermediate structure  $f'$ .

2. A statistical translation component 202 which translates  $f'$  into a corresponding intermediate target structure  $e'$ . This component incorporates a language model, a translation model, and a decoder.

3. A target transducer 203 which reconstructs a target sentence  $e$  from  $e'$ .

For statistical modeling, the target-translation transformation 203  $e' \rightarrow e$  is sometimes required to be invertible. Then  $e'$  can be constructed from  $e$  and no information is lost in the transformation.

The purpose of source and target transduction is to facilitate the task of the statistical translation. This will be accomplished if the probability distribution  $\Pr(f, e')$  is easier to model than the original distribution  $\Pr(f, e)$ . In practice this means that  $e'$  and  $f'$  should encode global linguistic facts about  $e$  and  $f$  in a local form.

A useful gauge of the success that a statistical model enjoys in modeling translation from sequences of source words represented by a random variable  $F$ , to sequences of target words represented by a random variable  $E$ , is the cross entropy<sup>3</sup>

<sup>3</sup>In this equation and in the remainder of this section, the convention is using uppercase letters (e.g.  $E$ ) for random variables and lower case letters (e.g.  $e$ ) for the values of random variables continues to be used.

$$H(E|F) = -\frac{1}{S} \sum_s \log P(e^s | f^s). \quad (165)$$

The cross entropy measures the average uncertainty that the model has about the target language translation  $e$  of a source language sequence  $f$ . Here  $P(e|f)$  is the probability according to the model that  $e$  is a translation of  $f$  and the sum runs over a collection of all  $S$  pairs of sentences in a large corpus comprised of pairs of sentences with each pair consisting of a source and target sentence which are translations of one another (See Sections 8–10).

A better model has less uncertainty and thus a lower cross entropy. The utility of source and target transductions can be measured in terms of this cross entropy. Thus transformations  $f \rightarrow f'$  and  $e' \rightarrow e$  are useful if models  $P'(f|e')$  and  $P'(e')$  can be constructed such that  $H(E'|F) < H(E|F)$ .

##### 11.4.2 Sense-Labeling in Statistical Translation

The remainder of this section is devoted to describing methods for constructing a sense-labeling transducer. In this case the following setup prevails:

The Intermediate Structures. Intermediate structures  $e'$  and  $f'$  consist of sequences of words labeled by their senses. Thus  $f'$  is a sentence over the expanded vocabulary whose

'words'  $f'$  are pairs  $(f, s)$  where  $f$  is a word in the original source vocabulary and  $s$  is its sense label. Similarly,  $e'$  is a sentence over the expanded vocabulary whose words  $e'$  are pairs  $(e, s)$  where  $e$  is a target word and  $s$  is its sense label.

Source and target transductions. For each source word and each target word, an informant site, such as first noun to the left is chosen, and an  $n$ -ary question about the value of the informant at that site is also chosen. A source-transduction transformation  $f \rightarrow f'$  and an inverse target-transduction transformation  $e' \rightarrow e$  map a sentence to the intermediate structure in which each word is labeled by a sense, determined by the question about its informant. A target-transduction transformation  $e' \rightarrow e$  maps a labeled sentence to a sentence in which the labels have been removed.

The probability models. A translation model such as one of the models in Sections 8–10 is used for both  $P'(F'|E')$  and for  $P(F|E)$ . A trigram language model such as that discussed in Section 6 is used for both  $P(E)$  and  $P(E')$ .

##### 11.4.3 The Viterbi Approximation

The probability  $P(f|e)$  computed by the translation model requires a sum over alignments as discussed in detail in Sections 8–10. This sum is often too expensive to compute directly since the number of alignments increases exponentially with sentence length. In the mathematical considerations of this Section, this sum will be approximated by the single term corresponding to the alignment,  $\mathcal{V}(f|e)$ , with greatest probability. This is the Viterbi approximation already discussed in Sections 8–10 and  $\mathcal{V}(f|e)$  is the Viterbi alignment.

Let  $c(f|e)$  be the expected number of time that  $e$  is aligned with  $f$  in the Viterbi alignment of a pair of sentences drawn at random from a large corpus of training data. Let  $c(\phi|e)$  be the expected number of times that  $e$  is aligned with  $\phi$  words. Then

$$c(f|e) = \frac{1}{S} \sum_s c(f|e; \mathcal{V}(f|e^s)) \quad (166)$$

$$c(\phi|e) = \frac{1}{S} \sum_s c(\phi|e; \mathcal{V}(\phi|e^s)).$$

where  $c(f|e; \mathcal{V})$  is the number of times that  $e$  is aligned with  $f$  in the alignment  $A$ , and  $c(\phi|e; \mathcal{V})$  is the number  $f$  times that  $e$  generates  $\phi$  target words in  $A$ . The counts above are also expressible as averages with respect to the model:

$$c(f|e) = \sum_{f, e} P(f, e) c(f|e; \mathcal{V}(f|e)) \quad (167)$$

$$c(\phi|e) = \sum_{f, e} P(f, e) c(\phi|e; \mathcal{V}(f|e)).$$

Probability distributions  $p(e, f)$  and  $p(\phi, e)$  are obtained by normalizing the counts  $c(f|e)$  and  $c(\phi|e)$ :

$$p(f, e) = \frac{1}{\text{norm}} c(f|e) p(\phi, e) = \frac{1}{\text{norm}} c(\phi, e). \quad (168)$$

<sup>4</sup>In these equations and in the remainder of the paper, the generic symbol

$$\frac{1}{\text{norm}}$$

is used to denote a normalizing factor that converts counts to probabilities. The actual value of

$$\frac{1}{\text{norm}}$$

will be implicit from the context. Thus, for example, in the left hand equation of (168), the normalizing factor is  $\text{norm} = \sum_{f,e} c(f|e)$  which equals the average length of source sentences. In the right hand equation of (168), the normalizing factor is the average length of target sentences.

(These are the probabilities that are stored in a table of probabilities 4205.) The conditional distributions  $p(f|e)$  and  $p(\phi|e)$  are the Viterbi approximation estimates for the parameters of the model. The marginals satisfy

$$\sum_{\phi} p(\phi, e) = u(e) \quad \sum_e p(f, e) = u(f) \quad (169)$$

$$\sum_f p(f, e) = \frac{1}{\text{norm}} \bar{\phi}(e) u(e)$$

where  $u(e)$  and  $u(f)$  are the unigram distributions of  $e$  and  $f$  and  $\bar{\phi}(e) = \sum_{\phi} p(\phi|e)$   $\phi$  is the average number of source words aligned with  $e$ . These formulae reflect the fact that in any alignment each source word is aligned with exactly one target word.

#### 11.4.4 Cross Entropy

In this subsection the cross entropies  $H(E|F)$  and  $H(E'|F)$  are expressed in terms of the information between source and target words.

In the Viterbi approximation, the cross entropy  $H(F|E)$  is given by,

$$H(F|E) = m \{ H(E|F) + H(\phi|E) \}, \quad (170)$$

where  $m$  is the average length of the source sentences in the training data, and  $H(F|E)$  and  $H(\phi|E)$  are the conditional entropies for the probability distributions  $p(f, e)$  and  $p(\phi, e)$ :

$$H(F|E) = - \sum_{f,e} p(f, e) \log p(f|e) \quad (171)$$

$$H(\phi|E) = - \sum_{e,\phi} p(\phi, e) \log p(\phi|e).$$

A similar expression for the cross entropy  $H(E|F)$  will now be given. Since

$$P(f, e) = P(f|e)P(e),$$

this cross entropy depends on both the translation model,  $P(f|e)$ , and the language model,  $P(e)$ . With a suitable additional approximation,

$$H(E|F) = m \{ H(\phi|E) - E(E, F) \} + H(E) \quad (172)$$

where  $H(E)$  is the cross entropy of  $P(E)$  and  $I(F, E)$  is the mutual information between  $f$  and  $e$  for the probability distribution  $p(f, e)$ .

The additional approximation required is,

$$H(F) = mH(F) = -m \sum_f p(f) \log p(f), \quad (173)$$

where  $p(f)$  is the marginal of  $p(f, e)$ . This amounts to approximating  $P(f)$  by the unigram distribution that is closest to it in cross entropy. Granting this, formula (172) is a consequence of (170) and of the identities

$$H(E|F) = H(E|F) - H(F) + H(E), \quad H(F) = H(F|E) + I(F, E). \quad (174)$$

Next consider  $H(E'|F)$ . Let  $e \rightarrow e'$  and  $f \rightarrow f'$  be sense labeling transformations of the type discussed above. Assume that these transformations preserve Viterbi alignments; that is, the words  $e$  and  $f$  are aligned in the Viterbi alignment for  $(f, e)$ , then their sensed versions  $e'$  and  $f'$  are aligned in the Viterbi alignment for  $(f', e')$ . It follows that the word translation probabilities obtained from the Viterbi alignments satisfy  $p(f, e) = \sum_{f'} p(f', e) = \sum_{e'} p(f, e')$ , where the sums range over the sensed versions  $f'$  of  $f$  and the sensed versions  $e'$  of  $e$ .

By applying (172) to the cross entropies  $H(E|F)$ ,  $H(E'|F)$ , and  $H(E'|F)$ , it is not hard to verify that

$$H(E|F) = H(E|F) - m \sum_f p(f) I(E, F|f), \quad (175)$$

$$H(E'|F) = H(E|F) - m \sum_e p(e) \{ I(F, E'|e) + I(\phi, E|e) \}.$$

Here  $I(E, F|f)$  is the conditional mutual information given a source word  $f$  between its translations  $E$  and its sensed versions  $F'$ ;  $I(F, E'|e)$  is the conditional mutual information given a target word  $e$  between its translations  $F$  and its sensed versions  $E'$ ; and  $I(\phi, E'|e)$  is the conditional mutual information given  $e$  between  $\phi$  and its sensed versions  $E'$ .

#### 11.5 Selecting Questions

The method depicted in 44 for finding good informants and questions for sensing is now described.

##### 11.5.1 Source Questions

For sensing source sentences, a question about an informant is a function  $\hat{c}$  from the source vocabulary into the set of possible senses. If the informant of  $f$  is  $x$ , then  $f$  is assigned the sense  $\hat{c}(x)$ . The function  $\hat{c}(x)$  is chosen to minimize the cross entropy  $H(E|F)$ . From formula (175), this is equivalent to maximizing the conditional mutual information  $I(F', E|f)$  between  $E$  and  $F'$

$$I(F', E|f) = \sum_{e,x} p(e, x|f) \log \frac{p(e, \hat{c}(x)|f)}{p(e|f)p(\hat{c}(x)|f)}, \quad (176)$$

where  $p(f, e, x)$  is the probability distribution obtained by counting the number of times in the Viterbi alignments that  $e$  is aligned with  $f$  and the value of the informant of  $f$  is  $x$ ,

$$p(f, e, x) = \frac{1}{\text{norm}} \sum_s c(f|e, x|f', e', s) \quad (177)$$

$$p(f, e) = \frac{1}{\text{norm}} \sum_{x: \hat{c}(x)=e} p(f, e, x).$$

An exhaustive search for the best  $\hat{c}$  requires a computation that is exponential in the number of values of  $x$  and is not practical. In the aforementioned paper entitled "Word-Sense Disambiguation using Statistical Methods" by the P. F. Brown, et al., a good  $\hat{c}$  is found using the flip-flop method which is only applicable if the number of senses is restricted to two. Here a different method that can be used to find  $\hat{c}$  for any number of senses is described. This method uses the technique of alternating minimization, and is similar to the k-means method for determining pattern clusters and to the generalized Lloyd method for designing vector quantizers.

The method is based on the fact that, up to a constant independent of  $\hat{e}$ , the mutual information  $I(F, E|f)$  can be expressed as an infimum over conditional probability distributions  $q(E|c)$ ,

$$I(F, E|f) = \inf_x \sum p(x) D(p(E|x, f); q(E|\hat{e}(x))) + \text{constant} \quad (178)$$

where

$$D(p(E); q(E)) = \sum_e p(e) \log \frac{p(e)}{q(e)}. \quad (179)$$

The best value of the information is thus an infimum over both the choice for  $\hat{e}$  and the choice for the  $q$ . This suggests the iterative method, depicted in 4401 for obtaining a good  $\hat{e}$ . This method comprises the steps of:

4401. Beginning with an initial choice of  $\hat{e}$ ;

4404. Performing Steps 4402 and 4403 until no further increase in  $I(F, E|f)$  results;

4403. For given  $q$ , finding the best  $\hat{e}$ :

$$\hat{e}(x) = \underset{e}{\operatorname{argmin}} D(p(E|x, f); q(E|e));$$

4402. For this  $\hat{e}$ , finding the best  $q$ :

$$q(e|c) = \frac{1}{\text{norm}} \sum_{x: \hat{e}(x)=e} p(e, x|f).$$

### 11.5.2 Target Questions

For sensing target sentences, a question about an informant is a function of  $\hat{e}$  from the target vocabulary into the set of possible senses.  $\hat{e}$  is chosen to minimize the entropy  $H(E|F)$ . From (175) this is equivalent to maximizing the sum

$$I(F, E|e) + I(\Phi, E|e).$$

In analogy to (179),

$$I(F, E|e) + I(\Phi, E|e) = \quad (180)$$

$$\inf_{q_1, q_2} \sum_x p(x) \{ D(p(F|x, e); q_1(F|\hat{e}(x))) + D(p(\Phi|x, e); q_2(\Phi|\hat{e}(x))) \}.$$

Again a good  $\hat{e}$  is obtained alternating minimization.

### 11.6 Generalizations

The methods of sense-labeling discussed above ask a single question about a single word of context. In other embodiments of the sense labeler, this question is the first question in a decision tree. In still other embodiments, rather than using a single informant site to determine the sense of a word, questions from several different informant sites are combined to determine the sense of a word. In one embodiment, this is done by assuming that the probability of an informant word  $x_i$  at informant site  $i$ , given a target word  $e$ , is independent of an informant word  $x_j$  at a different informant site  $j$  given the target word  $e$ . Also, in other embodiments, the intermediate source and target structure representations are more sophisticated than word sequences, including, but not limited to, sequences of lexical morphemes, case frame sequences, and parse tree structures.

### 12 ALIGNING SENTENCES

In this section, a method is described for aligning sentences in parallel corpora, and extracting from parallel corpora pairs of sentences which are translations of one another. These tasks are not trivial because at times a single sentence in one corpora is translated as two or more sentences in the other corpora. At other times a sentence, or even a whole passage, may be missing from one or the other of the corpora.

A number of researchers have developed methods that align sentences according to the words that they contain. (See for example, Deriving translation data from bilingual text, by R. Catizone, G. Russel, and S. Warwick, appearing in Proceedings of the First International Acquisition Workshop, Detroit, Mich. 1989; and "Making Connections", by M. Kay, appearing in ACH/ALLC '91, Tempe, Ariz. 1991.) Unfortunately, these methods are necessarily slow and, despite the potential for high accuracy, may be unsuitable for very large collections of text.

In contrast, the method described here makes no use of lexical details of the corpora. Rather, the only information that it uses, besides optional information concerning anchor points, is the lengths of the sentences of the corpora. As a result, the method is very fast and therefore practical for application to very large collections of text.

The method was used to align several million sentences from parallel French and English corpora derived from the proceedings of the Canadian Parliament. The accuracy of these alignments was in excess of 99% for a randomly selected set of 1000 alignments that were checked by hand. The correlation between the lengths of aligned sentences indicates that the method would achieve an accuracy of between 96% and 97% even without the benefit of anchor points. This suggests that the method is applicable to a very wide variety of parallel corpora for which anchor points are not available.

#### 12.1 Overview

One embodiment of the method is illustrated schematically in FIG. 46. It comprises the steps of:

4601 and 4602. Tokenizing the text of each corpus.

4603 and 4604. Determining sentence boundaries in each corpus.

4605. Determining alignments between the sentences of the two corpora.

The basic step 4605 of determining sentence alignments is elaborated further in FIG. 47. It comprises the steps of:

4701. Finding major and minor anchor points in each corpus. This divides each corpus into sections between major anchor points, and subsections between minor anchor points.

4702. Determining alignments between major anchor points.

4703. Retaining only those aligned sections for which the number of subsections is the same in both corpora.

4704. Determining alignments between sentences with each of the remaining aligned subsections.

One embodiment of the method will now be explained. The various steps will be illustrated using as an example the aforementioned parallel French and English corpora derived from the Canadian Parliamentary proceedings. These proceedings, called the Hansards, are transcribed in both English and French. The corpora of the example consists of the English and French Hansard transcripts for the years 1973 through 1986.

It is understood that the method and techniques illustrated in this embodiment and the Hansard example can easily be extended and adapted to other corpora and other languages.



### 12.2 Tokenization and Sentence Detection

First, the corpora are tokenized (steps 4601 and 4602) using a finite-state tokenizer of the sort described in Subsection 3.2.1. Next, (steps 4602 and 4603), the corpora are partitioned into sentences using a finite state sentence boundary detector. Such a sentence detector can easily be constructed by one skilled in the art. Generally, the sentences produced by such a sentence detector conform to the grade-school notion of sentence: they begin with a capital letter, contain a verb, and end with some type of sentence-final punctuation. Occasionally, they fall short of this ideal and consist merely of fragments and other groupings of words.

In the Hansard example, the English corpus contains 85,016,286 tokens in 3,510,744 sentences, and the French corpus contains 97,857,452 tokens in 3,690,425 sentences. The average English sentence has 24.2 tokens, while the average French sentence is about 9.5% longer with 26.5 tokens. The left-hand side of FIG. 48 shows the raw data for a portion of the English corpus, and the right-hand side shows the same portion after it was cleaned, tokenized, and divided into sentences. The sentence numbers do not advance regularly because the sample has been edited in order to display a variety of phenomena.

### 12.3 Selecting Anchor Points

The selection of suitable anchor points (Step 4071) is a corpus-specific task. Some corpora may not contain any reasonable anchors.

In the Hansard example, suitable anchors are supplied by various reference markers that appear in the transcripts. These include session numbers, names of speakers, time stamps, question numbers, and indications of the original language in which each speech was delivered. This auxiliary information is retained in the tokenized corpus in the form of comments sprinkled throughout the text. Each comment has the form \SCM{ } . . . \ECM{ } as shown on the right-hand side of FIG. 48.

To supplement the comments which appear explicitly in the transcripts, a number of additional comments were added. Paragraph comments were inserted as suggested by the space command of the original markup language. An example of this command appears in the eighth line of the left-hand side of FIG. 48. The beginning of a parliamentary session was marked by a Document comment, as illustrated in Sentence 1 on the right-hand side of FIG. 48. Usually, when a member addresses the parliament, his name is recorded. This was encoded as an Author comment, an example of which appears in Sentence 4. If the president speaks, he is referred to in the English corpus as Mr. Speaker and in the French corpus as

TABLE 18

Examples of comments	
English	French
Source = English	Source = Traduction
Source = Translation	Source = Francais
Source = Text	Source = Texte
Source = List Item	Source = List Item
Source = Question	Source = Question
Source = Answer	Source = Reponse

M. le Président. If several members speak at once, a shockingly regular occurrence, they are referred to as Some Hon. Members in the English and as Des Voix in the French. Times are recorded either as exact times on a 24-hour basis as in Sentence 81, or as inexact times of which there are two forms: Time=Later, and Time=Recess. These were encoded in the French as Time×Plus Tard and Time=Recess. Other types of comments are shown in Table 18.

The resulting comments laced throughout the text are used as anchor points for the alignment process. The comments Author=Mr. Speaker, Author=M. le Président, Author=Some Hon. Members, and Author=Des Voix are deemed minor anchors. All other comments are deemed major anchors with the exception of the Paragraph comment which was not treated as an anchor at all. The minor anchors are much more common than any particular major anchor, making an alignment based on minor anchors much less robust against deletions than one based on the major anchors.

### 12.4 Aligning Major Anchors

Major anchors, if they are to be useful, will usually appear in parallel in two corpora. Sometimes, however, through inattention on the part of translators or other misadventure, an anchor in one corpus may be garbled or omitted in another. In the Hansard example, for instance, this is problem is not uncommon for anchors based upon names of speakers.

The major anchors of two corpora are aligned (Step 4072) by the following method. First, each connection of an alignment is assigned a numeral cost that favors exact matches and penalizes omissions of garbled matches. In the Hansard example, these costs were chosen to be integers between 0 and 10. Connections between corresponding pairs such as Time=Later and Time=Plus Tard, were assigned a cost of 0, while connections between different pairs such as Time=Later and Author=Mr. Bateman were assigned a cost of 10. A deletion is assigned a cost of 5. A connection between two names were assigned a cost proportional to the minimal number of insertions, deletions, and substitutions necessary to transform one name, letter by letter, into the other.

Given these costs, the standard technique of dynamic programming is used to find the alignment between the major anchors with the least total cost. Dynamic programming is described by R. Bellman in the book titled Dynamic Programming, published by Princeton University Press, Princeton, N.J. in 1957. In theory, the time and space required to find this alignment grow as the product of the lengths of the two sequences to be aligned. In practice, however, by using thresholds and the partial traceback technique described by Brown, Spohrer, Hochschild, and Baker in their paper, Partial Traceback and Dynamic Programming, published in the Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing, in Paris, France in 1982, the time required can be made linear in the length of the sequences, and the space can be made constant. Even so, the computational demand is severe. In the Hansard example, the two corpora were out of alignment in places by as many as 90,000 sentences owing to mislabelled or missing files.

### 12.5 Discarding Sections

The alignment of major anchors partitions the corpora into a sequence of aligned sections. Next, (Step 4703), each section is accepted or rejected according to the population of minor anchors that it contains. Specifically, a section is accepted provided that, within the section, both corpora contain the same number of minor anchors in the same order. Otherwise, the section is rejected. Altogether, using this criteria, about 10% of each corpus was rejected. The minor anchors serve to divide the remaining sections into subsections that range in size from one sentence to several thousand sentences and average about ten sentences.

### 12.6 Aligning Sentences

The sentences within a subsection are aligned (Step 4074) using a simple statistical model for sentence lengths and paragraph markers. Each corpus is viewed as a sequence of sentence lengths punctuated by occasional paragraph markers, as illustrated in FIG. 49. In this figure, the circles around groups of sentence lengths indicate an alignment between the corresponding sentences. Each group is called a bead. The example consists of an ef-bead followed by an eff-bead followed by an e-bead followed by a  $\mathbb{L}_e\mathbb{L}_f$ -bead. From this perspective, an alignment is simply a sequence of beads that accounts for the observed sequences of sentence lengths and paragraph markers. The model assumes that the lengths of sentences have been generated by a pair of random processes, the first producing a sequence of beads and the second choosing the lengths of the sentences in each bead.

The length of a sentence can be expressed in terms of the number of tokens in the sentence, the number of characters in the sentence, or any other reasonable measure. In the Hansard example, lengths were measured as number of tokens.

The generation of beads is modelled by the two-state Markov model shown in FIG. 50. The allowed beads are shown in FIG. 51. A single sentence in one corpus is assumed to line up with zero, one, or two sentences in the other corpus. The probabilities of the different cases are assumed to satisfy  $\Pr(e)=\Pr(f)$ ,  $\Pr(ef)=\Pr(eef)$ , and  $\Pr(\mathbb{L}_e)=\Pr(\mathbb{L}_f)$ .

The generation of sentence lengths given beads is modeled as follows. The probability of an English sentence of length  $l_e$  given an e-bead is assumed to be the same as the probability of an English sentence of length  $l_e$  in the text as a whole. This probability is denoted by  $\Pr(l_e)$ . Similarly, the probability of a French sentence of length  $l_f$  given an f-bead is assumed to equal  $\Pr(l_f)$ . For an ef-bead, the probability of an English sentence of length  $l_e$  is assumed to equal  $\Pr(l_e)$  and the log of the ratio of length of the French sentence to the length of the English sentence is assumed to be normally distributed with mean  $\mu$  and variance  $\sigma^2$ . Thus, if  $r=\log(l_f/l_e)$ , then

$$\Pr(l_f/l_e)=\alpha \exp[-r(\mu)^2/2\sigma^2], \quad (181)$$

with  $\alpha$  chosen so that the sum of  $\Pr(l_f/l_e)$  over positive values of  $l_f$  is equal to unity. For an eef-bead, the English sentence lengths are assumed to be independent with equal marginals  $\Pr(l_e)$ , and the log of the ratio of the length of the French sentence to the sum of the lengths of the English sentences is assumed to be normally distributed with the same mean and variance as for an ef-bead. Finally, for an eff-bead, the probability of an English length  $l_e$  is assumed to equal  $\Pr(l_e)$  and the log of the ratio of the sum of the lengths of the French sentences to the length of the English sentence is assumed to be normally distributed as before. Then, given the sum of the lengths of the French sentences, the probability of a particular pair of lengths,  $l_{f1}$  and  $l_{f2}$ , is assumed to be proportional to  $\Pr(l_{f1})\Pr(l_{f2})$ .

Together, the model for sequences of beads and the model for sentence lengths given beads define a hidden Markov model for the generation of aligned pairs of sentence lengths. Markov Models are described by L. Baum in the article "An Inequality and associated maximization technique in statistical estimation of probabilistic functions of a Markov process", appearing in Inequalities in 1972.

The distributions  $\Pr(l_e)$  and  $\Pr(l_f)$  are determined from the relative frequencies of various sentence lengths in the data. For reasonably small lengths, the relative frequency is a reliable estimate of the corresponding probability. For longer lengths, probabilities are determined by fitting the observed frequencies of longer sentences to the tail of a

Poisson distribution. The values of the other parameters of the Markov model can be determined by from a large

TABLE 19

Parameter estimates	
Parameter	Estimate
$\Pr(e), \Pr(f)$	.007
$\Pr(ef)$	.690
$\Pr(eef), \Pr(eff)$	.020
$\Pr(\mathbb{L}_e), \Pr(\mathbb{L}_f)$	.005
$\Pr(\mathbb{L}_e\mathbb{L}_f)$	.245
$\mu$	.072
$\sigma^2$	.043

sample of text using the EM method. This method is described in the above referenced article by E. Baum.

For the Hansard example, histograms of the sentence length distributions  $\Pr(l_e)$  and  $\Pr(l_{f,en})$  for lengths up to 81 are shown in FIGS. 52 and 53 respectively. Except for lengths 2 and 4, which include a large number of formulaic sentences in both the French and the English, the distributions are very smooth.

The parameter values for the Hansard example are shown in Table 19. From these values it follows that 91% of the English sentences and 98% of the English paragraph markers line up one-to-one with their French counterparts. If  $X$  is a random variable whose log is normally distributed with mean  $\mu$  and variance  $\sigma^2$ , then the mean of  $X$  is  $\exp(\mu+\sigma^2/2)$ . Thus, from the values in the table, it also follows that the total length of the French text in an ef-, eef-, or eff-bead is about 9.8% greater on average than the total length of the corresponding English text. Since most sentences belong to ef-beads, this is close to the value of 9.5% given above for the amount by which the length of the average French sentences exceeds that of the average English sentence.

#### 12.7 Ignoring Anchors

For the Hansard example, the distribution of English sentence lengths shown in FIG. 52 can be combined with the conditional distribution of French sentence lengths given English sentence lengths from Equation (181) to obtain the joint distribution of French and English sentence lengths in ef-, eef-, and eff-beads. For this joint distribution, the mutual information between French and English sentence length is 1.85 bits per sentence. It follows that even in the absence of the anchor points, the correlation in sentence lengths is strong enough to allow alignment with an error rate that is asymptotically less than 100%.

Numerical estimates for the error rate as a function of the frequency of anchor points can be obtained by Monte Carlo simulation. The empirical distributions  $\Pr(l_e)$  and  $\Pr(l_f)$  shown in FIGS. 52 and 53, and the parameter values from Table 19 can be used to generated an artificial pair of aligned corpora, and then, the most probable alignment for these corpora can be found. The error rate can be estimated as the fraction of ef-beads in the most probable alignment that did not correspond to ef-beads in the true alignment.

By repeating this process many thousands of times, an expected error rate of about 0.9% was estimated for the actual frequency of anchor points in the Hansard data. By varying the parameters of the hidden Markov model, the effect of anchor points and paragraph markers on error rate can be explored. With paragraph markers but no anchor points, the expected error rate is 2.0%, with anchor points but no paragraph markers, the expected error rate is 2.3%, and with neither anchor points nor paragraph markers, the expected error rate is 3.2%. Thus, while anchor points and paragraph markers are important, alignment is still feasible

without them. This is promising since it suggests that the method is applicable to corpora for which frequent anchor points are not available.

TABLE 20

Unusual but correct alignments	
And love and kisses to you, too. ... mugwumps who sit on the fence with their mugs on one side and their wumps on the other side and do not know which side to come down on. At first reading, she may have.	Pareillement. ... en voulant ménager l'achève et le choux ils n'arrivent pas à prendre parti.  Elle semble en effet avoir un grief tout à fait valable, du moins au premier abord.

as data for construction of a statistical sense-labelling module.

### 12.8 Results for the Hansard Example

For the Hansard example, the alignment method described above ran for 10 days on an IBM Model 3090 mainframe under an operating system that permitted access to 16 megabytes of virtual memory. The most probable alignment contained 2,869,041 ef-beads. In a random sample of 1000 the aligned sentence pairs, 6 errors were found. This is consistent with the expected error rate of 0.9% mentioned above. In some cases, the method correctly aligned sentences with very different lengths. Examples are shown in Table 20.

### 13 ALIGNING BILINGUAL CORPORA

With the growing availability of machine-readable bilingual texts has become a burgeoning interest in methods for extracting linguistically valuable information from such texts. One way of obtaining such information is to construct sentence and word correspondences between the texts in the two languages of such corpora.

A method for doing this is depicted schematically in FIG. 45. This method comprises the steps of

4501. Beginning with a large bilingual corpus;

4504. Extracting pairs of sentences 4502 from this corpus such that each pair consists of a source and target sentence which are translations of each other;

4505. Within each sentence pair, aligning the words of the target sentence with the words in a source sentence, to obtain a bilingual corpus labelled with word-by-word correspondences 4503.

In one embodiment of Step 4504, pair of aligned sentences are extracted using the method explained in detail in Section 12. The method proceeds without inspecting the identities of the words within sentences, but rather uses only the number of words or numbers of characters that each sentence contains.

In one embodiment of Step 4505, word-by-word correspondences within a sentence pair are determined by finding the Viterbi alignment or approximate Viterbi alignment for the pair of sentences using a translation model of the sort discussed in Sections 8-10 above. These models constitute a mathematical embodiment of the powerfully compelling intuitive feeling that a word in one language can be translated into a word or phrase in another language.

Word-by-word alignments obtained in this way offer a valuable resource for work in bilingual lexicography and machine translation. For example, a method of cross-lingual sense labeling, described in Section 11, and also in the aforementioned paper, "Word Sense Disambiguation using Statistical Methods", uses alignments obtained in this way

### 14 Hypothesis Search - Steps 702 and 902

#### 14.2 Overview of Hypothesis Search

Referring now to FIG. 7, the second step 702 produces a set of hypothesized target structures which correspond to putative translations of the input intermediate source structure produced by step 701. The process by which these target structures are produced is referred to as hypothesis search. In a preferred embodiment target structures correspond to sequences of morphemes. In other embodiments more sophisticated linguistic structures such as parse trees or case frames may be hypothesized.

An hypothesis in this step 702 is comprised of a target structure and an alignment of a target structure with the input source structure. Associated with each hypothesis is a score. In other embodiments a hypothesis may have multiple alignments. In embodiments in which step 701 produces multiple source structures an hypothesis may contain multiple alignments for each source structure. It will be assumed here that the target structure comprised by a hypothesis contains a single instance of the null target morpheme. The null morphemes will not be shown in the Figures pertaining to hypothesis search, but should be understood to be part of the target structures nonetheless. Throughout this section on hypothesis search, partial hypothesis will be used interchangeably with hypothesis, partial alignment with alignment, and partial target structure with target structure.

The target structures generated in this step 702 are produced incrementally. The process by which this is done is depicted in FIG. 54. This process is comprised in five steps.

A set of partial hypotheses is initialized in step 5401. A partial hypothesis is comprised of a target structure and an alignment with some subset of the morphemes in the source structure to be translated. The initial set generated by step 5401 consists of a single partial hypothesis. The partial target structure for this partial hypothesis is just an empty sequence of morphemes. The alignment is the empty alignment in which no morphemes in the source structure to be translated are accounted for.

The system then enters a loop through steps 5402, 5403, and 5404, in which partial hypotheses are iteratively extended until a test for completion is satisfied in step 5403. At the beginning of this loop, in step 5402, the existing set of partial hypotheses is examined and a subset of these hypotheses is selected to be extended in the steps which comprise the remainder of the loop. In step 5402 the score for each partial hypothesis is compared to a threshold (the method used to compute these thresholds is described below). Those partial hypotheses with scores greater than threshold are then placed on a list of partial hypotheses to be extended in step 5404. Each partial hypothesis that is

extended in step 5404 contains an alignment which accounts for a subset of the morphemes in the source sentence. The remainder of the morphemes must still be accounted for. Each extension of an hypothesis in step 5404 accounts for one additional morpheme. Typically, there are many tens or hundreds of extensions considered for each partial hypothesis to be extended. For each extension a new score is computed. This score contains a contribution from the language model as well as a contribution from the translation model. The language model score is a measure of the plausibility a priori of the target structure associated with the extension. The translation model score is a measure of the plausibility of the partial alignment associated with the extension. A partial hypothesis is considered to be a full hypothesis when it accounts for the entire source structure to be translated. A full hypothesis contains an alignment in which every morpheme in the source structure is aligned with a morpheme in the hypothesized target structure. The iterative process of extending partial hypotheses terminates when step 5402 produces an empty list of hypotheses to be extended. A test for this situation is made on step 5403.

This method for generating target structure hypotheses can be extended to an embodiment of step 902 of FIG. 9, by modifying the hypothesis extension step 5404 in FIG. 54, with a very similar step that only considers extensions which are consistent with the set of constraints 906. Such a modification is a simple matter for one skilled in the art.

#### 14.2 Hypothesis Extension 5404

This section provides a description of the method by which hypotheses are extended in step 5404 of FIG. 54. Examples will be taken from an embodiment in which the source language is French and the target language is English. It should be understood however that the method described applies to other language pairs.

##### 14.2.1 Types of Hypothesis Extension

There are a number of different ways a partial hypothesis may be extended. Each type of the method described in this section assigns scores to partial hypotheses based on Model 3 from the section entitled Translation Models and Parameter Estimation. One skilled in the art will be able to adopt the method described here to other modes, such as Model 5, which is used in the best mode of operation.

A partial hypothesis is extended by accounting for one additional previously unaccounted for element of the source structure. When a partial hypothesis  $H_1$  is extended to some other hypothesis  $H_2$ , the score assigned to  $H_2$  is a product of the score associated with  $H_1$  and a quantity denoted as the extension score. The value of the extension score is determined by the language mode, the translation model, the hypothesis being extended and the particular extension that is made. A number of different types of extensions are possible and are scored differently. The possible extension types and the manner in which they are scored is illustrated in the examples below.

As depicted in FIG. 55, in a preferred embodiment, the French sentence *La jeune fille a reveillé sa mère* is transduced in either step 701 of FIG. 7 or step 901 of FIG. 9 into the morphological sequence *la jeune fille V<sub>-part</sub> 3s sa mère*.

The initial hypothesis accounts for no French morphemes and the score of this hypothesis is set to 1. This hypothesis can be extended in a number of ways. Two sample extensions are shown in FIGS. 56 and 57. In the first example in FIG. 56, the English morpheme *the* is hypothesized as accounting for the French morpheme *la*. The component of the score associated with this extension is the equal to

$$l(\text{the}|*,*)n(1|\text{the})t(\text{la}|\text{the})d(1|1).$$

(182)

Here,  $*$  is a symbol which denotes a sequence boundary, and the factor  $l(\text{the}|*,*)$  is the trigram language model parameter that serves as an estimate of the probability that the English morpheme *the* occurs at the beginning of a sentence. The factor  $n(1|\text{the})$  is the translation model parameter that is an estimate of the probability that the English morpheme *the* has fertility 1, in other words, that the English morpheme *the* is aligned with only a single French morpheme. The factor  $t(\text{la}|\text{the})$  is the translation model parameter that serves as an estimate of the lexical probability that the English morpheme *the* translates to the French morpheme *la*. Finally, the factor  $d(1|1)$  is the translation model parameter that serves as an estimate of the distortion probability that a French morpheme will be placed in position 1 of the French structure given that it is aligned with an English morpheme that is in position 1 of the English structure. In the second example in FIG. 57, the English morpheme *mother* is hypothesized as accounting for the French morpheme *mère*. The score for this partial hypothesis is

$$l(\text{mother}|*,*)n(1|\text{mother})t(\text{mère}|\text{mother})d(7|1).$$

(183)

Here, the final  $d(7|1)$  serves as an estimate of the distortion probability that a French morpheme, such as *mère*, will be placed in the 7th position in a source sequence given that it is aligned with an English morpheme such as *mother* which is in the 1st position in an hypothesized target sequence.

Now, suppose the partial hypothesis in FIG. 56 is to be extended on some other invocation of step 5404. A common translation of the pair of French morphemes *jeune fille* is the English morpheme *girl*. However, since in a preferred embodiment a partial hypothesis is extended to account for only a single French morpheme at a time, it is not possible to account for both *jeune* and *fille* with a single extension. Rather the system first accounts for one of the morphemes, and then on another round of extensions, accounts for the other. This can be done in two ways, either by account first for *jeune* or by accounting first for *fille*. FIG. 58 depicts the extension that accounts first for *fille*. The  $+$  symbol in FIG. 58 after the the English morpheme *girl* denotes the fact that in these extensions *girl* is to be aligned with more French morphemes than it is currently aligned with, in this case, at least two. A morpheme so marked is referred to as open. A morpheme that is not open is said to be closed. A partial hypothesis which contains an open target morpheme is referred to as open, or as an open partial hypothesis. A partial hypothesis which is not open is referred to as closed, or as a closed partial hypothesis. An extension is referred to as either open or closed according to whether or not the resultant partial hypothesis is open or closed. In a preferred embodiment, only the last morpheme in a partial hypothesis can be designated open. The score for the extension in FIG. 58 is

$$l(\text{girl}|*,\text{be})2\left(\sum_{i=2}^{25} n(i|\text{girl})\right)t(\text{fille}|\text{girl})d(3|2).$$

(184)

Here, the factor  $l(\text{girl}|*,\text{be})$  is the language model parameter that serves as an estimate of the probability with which the English morpheme *girl* is the second morpheme in a source structure in which the first morpheme is *the*. The next factor of 2 is the combinatorial factor that is discussed in the section entitled Translation Models and Parameter Estimation. It is factored in, in this case, because the open English

morpheme girl is to be aligned with at least two French morphemes. The factor  $n(i|girl)$  is the translation model parameter that serves as an estimate of the probability that the English morpheme girl will be aligned with exactly  $i$  French morphemes, and the sum of these parameters for  $i$  between 2 and 25 is an estimate of the probability that girl will be aligned with at least 2 morphemes. It is assumed that the probability that an English morpheme will be aligned with more than 25 French morphemes is 0. Note that in a preferred embodiment of the present invention, this sum can be prepared and stored in memory as a separate parameter. The factor  $t(fille|girl)$  is the translation model parameter that serves as an estimate of the lexical probability that one of the French morphemes aligned with the English morpheme girl will be the French morpheme fille. Finally, the factor  $d(3|2)$  is the translation model parameter that serves as an estimate of the distortion probability that a French morpheme will be placed in position 3 of the French structure given that is aligned with an English morpheme which is in position 2 of the English structure. This extension score in Equation 184 is multiplied by the score in Equation 182 for the partial hypothesis which is being extended to yield a new score for the partial hypothesis in FIG. 56 of

$$l(the|*,*)n(1|the)l(the)d(1|1) \times \quad (185)$$

$$l(girl|*,the)2 \left( \frac{25}{\sum_{i=2}^{25} n(i|girl)} \right) t(fille|girl)d(3|2). \quad (186)$$

Consider now an extension to the partial hypothesis in FIG. 58. If a partial hypothesis that is to be extended contains an open morpheme, then, in a preferred embodiment, that hypothesis can only be extended by aligning another morpheme from the source structure with that open morpheme. When such an extension is made, there are two possibilities: 1) the open morpheme is kept open in the extended partial hypothesis, indicating that more source morphemes are to be aligned with that open target morpheme, or 2) the open morpheme is closed indicating that no more source morphemes are to be aligned with that target morpheme. These two cases are illustrated in FIGS. 59 and 60.

In FIG. 59, an extension is made of the partial alignment in FIG. 58 by aligning the additional French morpheme jeune with the English morpheme girl. In this example the English morpheme girl is then closed in the resultant partial hypothesis. The extension score for this example is

$$\left( \frac{n(2|girl)}{\left( \frac{25}{\sum_{i=2}^{25} n(i|girl)} \right)} \right) t(jeune|girl)d(2|2).$$

Here, the first quotient adjusts the fertility score for the partial hypothesis by dividing out the estimate of the probability that girl is aligned with at least two French morphemes and by multiplying in an estimate of the probability that girl is aligned with exactly two French morphemes. As in the other examples, the second and third factors are estimates of the lexical and distortion probabilities associated with this extension.

In FIG. 60, the same extension is made as in FIG. 59, except here the English morpheme girl is kept open after the extension, hence the + sign. The extension score for this example is

$$3 \left( \frac{\left( \frac{25}{\sum_{i=3}^{25} n(i|girl)} \right)}{\left( \frac{25}{\sum_{i=2}^{25} n(i|girl)} \right)} \right) t(jeune|girl)d(2|2). \quad (187)$$

Here, the factor of 3 is the adjustment to the combinatorial factor for the partial hypothesis. Since the score for the partial hypothesis in FIG. 59 already has a combinatorial factor of 2, the score for the resultant partial hypothesis in FIG. 60, will have a combinatorial factor of  $2 \times 3 = 3!$ . The quotient adjust the fertility score for the partial hypothesis to reflect the fact that in further extensions of this hypothesis girl will be aligned with at least three French morphemes.

Another type of extension performed in the hypothesis search is one in which two additional target morphemes are appended to the partial target structure of the partial hypothesis being extended. In this type of extension, the first of these two morphemes is assigned a fertility of zero and the second is aligned with a single morpheme from the source structure. This second target morpheme may be either open or closed.

FIG. 62 shows an extension of the partial hypothesis in FIG. 61 by the two target morphemes up her, in which her is aligned with the source morpheme sa. The score for this extension is

$$l(up|girl, to\_wake)l(her|to\_wake, up)n(0|up)n(1|her)t(sa|her)d(6|6). \quad (188)$$

Here, the first two factors are the trigram language model estimates of the probabilities with which up follows girl to\\_wake, and with which her follows to\\_wake up, respectively. The third factor is the fertility parameter that serves as an estimate of the probability that up is aligned with no source morphemes. The fourth, fifth, and sixth factors are the appropriate fertility, lexical, and distortion parameters associated with the target morpheme her in this partial alignment.

FIG. 63 shows a similar extension by up her. The difference with the extension in FIG. 62 is that in FIG. 63, the source morpheme her is open. The score for this extension is

$$l(up|girl, to\_wake)l(her|to\_wake, up)n(0|up) \times \quad (189)$$

$$2 \left( \frac{25}{\sum_{i=2}^{25} n(i|her)} \right) t(sa|her)d(6|6).$$

The score for this extension differs from the score in Equation 188 in that the fertility parameter  $n(1|her)$  is replaced by the combinatorial factor 2 and the sum of fertility parameters which provides an estimate of the probability that her will be aligned with at least two source morphemes.

A remaining type of extension is where a partial hypothesis is extended by an additional connection which aligns a source morpheme with the null target morpheme. The score for this type of extension is similar to those described above. No language model score is factored in, and scores from the translation model are factored in, in accordance with the

probabilities associated with the null word as described in the section entitled Translation Models and Parameter Estimation.

#### 14.3 Selection of Hypothesis to Extend 5402

Throughout the hypothesis search process, partial hypotheses are maintained in a set of priority queues. In theory, there is a single priority queue for each subset of positions in the source structure. So, for example, for the source structure *oui, oui, three positions, oui* is in position 1; a comma is in position 2; and *oui* is in position 3, and there are therefore  $2^3$  subsets of positions: [], [1], [2], [3], [1,2], [1,3], [2,3], and [1,2,3]. In practice these priority queues are initialized only on demand, and many less than the full number of queues possible are used in the hypothesis search. In a preferred embodiment, each partial hypothesis is comprised of a sequence of target morphemes, and these morphemes are aligned with a subset of source morphemes. Corresponding to that subset of source morphemes is a priority queue in which the partial hypothesis is stored. The partial hypotheses within a queue are prioritized according to the scores associated with those hypotheses. In certain preferred embodiments the priority queues are limited in size and only the 1000 hypothesis with the best scores are maintained.

The set of all subsets of a set of source structure positions can be arranged in a subset lattice. For example, the subset lattice for the set of all sets of the set [1,2,3] is shown in FIG. 64. In a subset lattice, a parent of a set *S* is any which contains one less element than *S*, and which is also a subset of *S*. In FIG. 64 arrows have been drawn from each set in the subset lattice to each of its parents. For example, the set [2] is a parent of the set [1,2].

A subset lattice defines a natural partial ordering on a set of sets. Since the priority queues used in hypothesis search are associated with subsets, a subset lattice also defines a natural partial ordering on the set of priority queues. Thus in FIG. 64, there are two parents of the priority queue associated with the subset of source structure positions [1,3]. These two parents are the priority queues associated with the set [1] and [3]. A priority queue  $Q_1$  is said to be an ancestor of another priority queue  $Q_2$  if 1)  $Q_1$  is not equal to  $Q_2$ , and 2)  $Q_1$  is a subset of  $Q_2$ . If  $Q_1$  is an ancestor of  $Q_2$ , then  $Q_2$  is said to be a descendant of  $Q_1$ .

Considering now the process by which a set of partial hypotheses are selected in step 5402 to be extended in step 5404, when step 5402 is invoked, it is invoked with a list of partial hypotheses that were either 1) created by the initialization step 5401, or 2) created as the result of extensions in step 5404 on a previous pass through the loop comprised of steps 5402, 5403, and 5404. These partial hypotheses are stored in priority queues according to the sets of source morphemes they account for. For example, the partial hypothesis in FIG. 65 would be stored in the priority queue associated with the set [2,3], since it accounts for the source morphemes in positions 2 and 3.

A priority queue is said to be active if there are partial hypotheses stored in it. An active priority queue is said to be on the frontier if it has no active descendant. The cardinality of a priority queue is equal to the number of elements in the subset with which it is associated. So, for example, the cardinality of the priority queue which is associated with the set [2,3] is 2.

The process in step 5402 functions by assigning a threshold to every active priority queue and then places on the list of partial hypotheses to be extended every partial hypothesis on an active priority queue that has a score that is greater than the threshold for that priority queue. This is depicted in FIG. 66. First, in step 6601 the threshold for every active

priority queue is initialized to infinity, in practice, some very large number. Second, in step 6602, thresholds are determined for every priority queue on the frontier.

The method by which these thresholds are computed is best described by first describing what the normalizer of a priority queue is. Each priority queue on the frontier corresponds to a set of positions of source morphemes. At each position of these positions is a particular source morpheme. Associated with each morpheme is a number, which in a preferred embodiment is the unigram probability of that source morpheme. These unigram probabilities are estimated by transducing a large body of source text and simply counting the frequency with which the different source morphemes occur. The normalizer for a priority queue is defined to be the product of all the unigram probabilities for the morphemes at the positions in the associated set of source structure positions. For example, the normalizer for the priority queue associated with the set [2,3] for the source structure *la jeune fille V\_past\_3s réveiller sa mère* is:

$$\text{normalizer}([2,3]) = \text{Pr}(\text{jeune})\text{Pr}(\text{fille}). \quad (190)$$

For each priority queue *Q* on the frontier define the normed score of *Q* to be equal to the score of the partial hypothesis with the greatest score in *Q* divided by the normalizer for *Q*. Let *Z* be equal to the maximum of all normed scores for all priority queues on the frontier. The threshold assigned to a priority queue *Q* on the frontier is then equal to *Z* times the normalizer for that priority queue divided by a constant which in a preferred embodiment is 45.

After step 6602, thresholds have been assigned to the priority queues on the frontier, a loop is performed in steps 6604 through 6610. The loop count *i* is equal to a different cardinality on each iteration of the loop. The counter *i* is initialized in step 6604 to the largest cardinality of any active priority queue, in other words, *i* is initialized to the maximum cardinality of any priority queue on the frontier. On each iteration of the loop the value of *i* is decreased by 1 until *i* is equal to 0, at which point the test 6604 is satisfied and the process of selecting partial hypotheses to be extended is terminated.

Inside the loop through cardinalities is another loop in steps 6606 through 6609. This is a loop through all active priority queues of a given cardinality. In this loop each priority queue of cardinality *i* is processed in step 6608.

A schematic flow diagram for this processing step 6608 is shown in FIG. 67. The priority queue *Q* to be processed enters this step at 6701. Steps 6704 through 6707 perform a loop through all partial hypotheses *i* in the priority queue *Q* which are greater than the threshold associated with *Q*. At step 6705 the partial hypothesis *i* is added to the list of partial hypotheses to be extended. At step 6706 *i* is used to adjust the thresholds of all active priority queues which are parents of *Q*. These thresholds are then used when priority queues of lower priority are processed in the loop beginning at step 6604 in FIG. 66.

Each priority queue which is a parent of partial hypothesis *i* at step 6706 contains partial hypotheses which account for one less source morpheme than the partial hypothesis *i* does. For example, consider the partial hypothesis depicted in FIG. 59. Suppose this is the partial hypothesis *i*. The two target morphemes *the* and *girl* are aligned with the three source morphemes *la, jeune, and fille* which are in source structure positions 1, 2, and 3 respectively. This hypothesis *i* is therefore in the priority queue corresponding to the set [1,2,3]. The priority queues that are parents of this hypothesis correspond to the sets [1,2], [1,3], and [2,3]. We can use



partial hypothesis *i* to adjust the threshold in each of these priority queues, assuming they are all active, by computing a parent score,  $score_p$ , from the score  $score_i$  associated with the partial hypothesis *i*. A potentially different parent score is computed for each active parent priority queue. That parent score is then divided by a constant, which in a preferred embodiment is equal to 45. The new threshold for that queue is then set to the minimum of the previous threshold and that parent score.

These parent scores are computed by removing from  $score_i$  the contributions for each of the source morphemes *la*, *jeune*, and *fille*. For example, to adjust the threshold for the priority queue [2,3], it is necessary to remove the contribution to the score, associated with the source morpheme in position 1, which is *la*. This morpheme is the only morpheme aligned with the, so the language model contribution for the must be removed, as well as the translation model contributions associated with *la*. Therefore:

$$score_p = \frac{score_i}{((the) * boundary) * ((la) * the) * n(1) * the * d(1,1)} \quad (191)$$

As another example, to adjust the threshold for the priority queue [1,3], it is necessary to remove the contribution to the score, associated with the source morpheme in position 2, which is *jeune*. This morpheme is one of two aligned with the target morpheme *girl*. If the connection between *girl* and *jeune* is removed from the partial alignment in FIG. 59, there is still a connection between *girl* and *fille*. In other words, *girl* is still needed in the partial hypothesis to account for *fille*. Therefore, no language model component is removed. The parent score in this case is:

$$score_p = score_i \cdot \frac{\sum_{j=2}^{25} n(j|girl)}{n(2|girl)} \cdot \frac{1}{n(jeune|girl)} \cdot \frac{1}{d(2,2)} \quad (192)$$

Here, the first quotient adjust the fertility score, the second adjusts the lexical score and the third adjusts the distortion score.

With some thought, it will be clear to one skilled in the art how to generalize from these examples to other situations. In general, a parent score is computed by removing a connection from the partial alignment associated with the partial hypothesis *i*. Such a connection connects a target morpheme *t* in the partial target structure associated with the partial hypothesis *i* and a source morpheme *s* in a source structure. If this connection is the only connection to the target morpheme *t*, then the language model score for *t* is divided out, otherwise it is left in. The lexical and distortion scores associated with the source morpheme *s* are always divided out, as is the fertility score associated with the target morpheme *t*. If *n* connections remain to the target morpheme *t*, since *n*+1 source morphemes are aligned with *t* in the partial hypothesis *i*, then the open fertility score serving as an estimate of the probability that at least *n*+1 source morphemes will be aligned with *t* is multiplied in.

When the list of hypotheses to be extended that is created in step 5402 is empty the search terminates.

Refer now to step 5404 in FIG. 54. This step extends a list of partial hypotheses. An embodiment of the method by which this extension takes place is documented in the pseudo code in FIG. 68.

The procedure `extend_partial_hypotheses_on_list` takes as input a list of hypotheses to be extended. Lines 1 through 8 contain a loop through all the partial hypotheses on the list which are extended in turn. In line 3 the variable *h* is set to the partial hypothesis being extended on iteration of the loop. For each hypothesis, *h*, that is extended, it can

be extended by aligning an additional source morpheme with a morpheme in an hypothesized target structure. In lines 4 through 7 a loop is made through every position *p* in the source structure. In certain embodiments, a loop may be made only through the first *n* source positions that are not already aligned in the partial hypothesis *h*. In a preferred embodiment *n* is set to 4. At line 5 a test is made to see if the source morpheme at position *p* is already aligned by the partial hypothesis with an hypothesized target morpheme. If it is not, then at line 6 a call is made to the procedure `extend_h_by_accounting_for_source_morpheme_in_position_p` is contained in lines 10 through 31 of Figure. At line 11 a check is made to determine if the partial hypothesis *h* is open, in other words, if it contains an open target morpheme. If it is open then extensions are made in lines 12 through 14. On line 12, the variable *q* is set to the position in the hypothesized partial target structure of the open morpheme. Each of these extensions made in lines 12 through 14 are made by adding a connection to the partial alignment of *h*. Each such connection is a connection from the morpheme at position *p* in the source structure to the open morpheme at position *q* in the target structure. On line 13, an extension is created in which a connection from *p* to *q* is added to the partial alignment of *h* and in which the morpheme at position *t* is kept open. On line 14, an extension is created in which a connection from *p* to *q* is added to the partial alignment of *h* and in which the morpheme at position *t* is closed.

Extensions of partial hypotheses *h* which are closed are made in lines 17 through 29. First, in line 17 the variable *s* is set to the identity of the source morpheme at position *p* in the source structure. This morpheme will have a number of possible target translations. In terms of the translation model, this means that there will be a number of target morphemes *t* for which the lexical parameter  $t(i|s)$  is greater than a certain threshold, which in an embodiment is set equal to 0.001. The list of such target morphemes for a given source morpheme *s* can be precomputed. In lines 18 through 29 a loop is made through a list of the target morphemes for the source morpheme *s*. The variable *t* is set to the target morpheme being processed in the loop. On line 20, an extension is made in which the target morpheme *t* is appended to the right end of the partial target structure associated with *h* and then aligned with the source morpheme at position *p*, and in which the target morpheme *t* is open in the resultant partial hypothesis. On line 21, an extension is made in which the target morpheme *t* is appended to the right end of the partial target structure associated with *h* and then aligned with the source morpheme at position *p*, and in which the target morpheme *t* is closed in the resultant partial hypothesis. On line 22, an extension is made in which the target morpheme *t* is appended to the null target morpheme in the partial target structure associated with hypothesis *h*. It is assumed throughout this description of hypothesis search that every partial hypothesis comprises a single null target morpheme.

The remaining types of extensions to be performed are those in which the target structure is extended by two morphemes. In such extensions, the source morpheme at position *p* is aligned with the second of these two target morphemes. On line 23, a procedure is called which creates a list of target morphemes that can be inserted between the last morpheme on the right of the hypothesis *h* and the hypothesized target morpheme, *t*. The lists of target morphemes created by this procedure can be precomputed from language model parameters. In particular, suppose  $t_p$  is the last morpheme on the right of the partial target structure

comprised by the partial hypothesis  $h$ . For any target morpheme  $t_i$ , the language model provides a score for the three-word sequence  $t_i, t_i, t_i$ . In one preferred embodiment this score is equal to an estimate of 1-gram probability for the morpheme  $t_i$ , multiplied by an estimate of the probability with 2-gram conditional probability with which  $t_i$  follows  $t_i$ , multiplied by an estimate of the 3-gram conditional probability with which  $t_i$  follows the pair  $t_i, t_i$ . By computing such a score for each target morpheme  $t_i$ , the target morphemes can be ordered according to these scores. The list returned by the procedure called on line 23 is comprised of the  $m$  best  $t_i$ 's which have scores greater than a threshold  $z$ . In one embodiment,  $z$  is equal to 0.001 and  $m$  is equal to 100.

The loop on lines 24 through 28 makes extensions for each  $t_i$  on the list created on line 23. On lines 26 and 27, extensions are made in which the pair of target morphemes  $t_i, t_i$  is appended to the end of the partial target structure comprised by the partial hypothesis  $h$ , and in which the source morpheme at position  $p$  is aligned with  $t_i$ . The hypothesis which result from extensions made on line 23 are open and the hypotheses which result from extensions made on line 23 are closed.

While various embodiments of the present invention have been described above, it should be understood that they have been presented by way of example, and not limitation. Thus the breadth and scope of the present invention should not be limited by any of the above-described exemplary embodiments, but should be defined only in accordance with the following claims and their equivalents.

Having thus described our invention, what we claim as new and desire to secure by Letters Patent is:

1. A method operating on a computer for translating source text from a first language to target text in a second language different from the first language, said method comprising the steps of:

measuring a value of the source text in the first language and storing the source text in a first memory buffer;

generating the target text in the second language based on a combination of a probability of occurrence of an intermediate structure of text associated with a target hypothesis selected from the second language using a target language model, and a probability of occurrence of the source text given the occurrence of said intermediate structure of text associated with said target hypothesis using a target-to-source translation model; and

performing at least one of a storing operation to save said target text in a second memory buffer and a presenting operation to make said target text available for at least one of a viewing and listening operation using an I/O device.

2. A method according to claim 1, further comprising the step of:

receiving at least one user defined criteria pertaining to the source text to thereby bound the target text, said criteria belonging to the group comprising specification of a noun phrase, specification of subject, specification of verb, a semantic sense of a word or group of words, inclusion of a predetermined word or group of words, and exclusion of a predetermined word or group of words.

3. A method operating on a computer for translating source text from a first language to target in a second language different from the first language, said method comprising the steps of:

receiving the source text in the first language;

generating at least one target hypothesis, each of said target hypotheses comprising text selected from the second language;

estimating, for each target hypothesis, a first probability of occurrence of said text associated with said target hypothesis using a target language model;

estimating, for each target hypothesis, a second probability of occurrence of the source text given the occurrence of said text associated with said target hypothesis using a target-to-source translation model;

combining, for each target hypothesis, said first and second probabilities to produce a target hypothesis match score; and

performing at least one of a storing operation to save and a presenting operation to make available for at least one of a viewing and listening operation, at least one of said target hypotheses according to its associated match score.

4. A method according to claim 3, further comprising the steps of:

generating a list of partial target hypotheses using said target language model and said target-to-source translation model;

extending said partial hypotheses to generate a set of new hypotheses;

determining a match score for at least one of said new hypotheses; and

performing at least one of a storing operation to save and a presenting operation to make available for viewing, at least one of said new hypotheses according to its associated match score.

5. A method according to claim 4, wherein said match score determining step further comprises the steps of:

determining a language model score for each of said partial hypotheses using said target language model;

determining a translation model score for each of said partial hypotheses using said target-to-source translation model; and

determining a combined score of said language and translation model scores for each of said partial hypotheses.

6. A method according to claim 5, wherein said translation model score determining step further comprises at least one of the following steps:

determining a fertility model score for each of said partial hypotheses, said partial hypotheses comprising at least one notion and said fertility model score being proportional to a probability that a notion in the target text will generate a specific number of units of linguistic structure in the source text;

determining an alignment score for each of said partial hypotheses, said alignment score being proportional to a probability that a unit of linguistic structure in the target text will align with one of zero or more units of linguistic structure in the source text;

determining a lexical model score, for each of said partial hypotheses, said lexical model score being proportional to the probability that said units of linguistic structure in the target text of a given partial hypothesis will translate into said units of linguistic structure of said source text;

determining a distortion model score for each of said partial hypotheses, said distortion model score being proportional to the probability that a source unit of



101

linguistic structure will be in a particular position given a position of the target units of linguistic structure that generated it; and

determining a combined score, said combined score being proportional to a product of at least one of said fertility, alignment, lexical and distortion model scores for each of said partial hypotheses.

7. A method according to claim 5, wherein said translation model score determining step further comprises at least one of the following steps:

determining a fertility model score for each of said partial hypotheses, said partial hypotheses comprising at least one notion and said fertility model score being proportional to a probability that a notion in the target text will generate a specific number of notion units of linguistic structure in an intermediate structure source text;

determining an alignment score for each of said partial hypotheses, said alignment score being proportional to the probability that a unit of linguistic structure in said intermediate structure of target text will align with one of zero or more units of linguistic structure in said intermediate structure of source text;

determining a lexical model score, for each of said partial hypotheses, said lexical model score being proportional to a probability that said units of linguistic structure in said intermediate structure of target text of a given partial hypothesis will translate into said units of linguistic structure of said intermediate structure source text;

determining a distortion model score for each of said partial hypotheses, said distortion model score being proportional to a probability that a source unit of linguistic structure will be in a particular position given the position of the target units of linguistic structure that generated it; and

determining a combined score, said combined score being proportional to a product of at least one of said fertility, alignment, lexical and distortion model scores for each of said partial hypotheses.

8. A method according to claim 6, wherein each of said notions comprises at least one unit of linguistic structure, and each of said units of linguistic structure in the target text produces one of zero or more units of linguistic structure in the source text.

9. A method according to claim 3, further comprising the step of:

receiving at least one user defined criteria pertaining to the source text to thereby bound the target text, criteria belonging to the group comprising specification of a noun phrase, specification of subject, specification of verb, a semantic sense of a word or group of words, inclusion of a predetermined word or group of words, and exclusion of a predetermined word or group of words.

10. A method operating on a computer for translating source text from a first language to target text in a second language different from the first language, said method comprising the steps of:

receiving the source text in the first language and storing the source text in a first memory buffer;

receiving one of zero or more user defined criteria pertaining to the source and target texts to thereby bound the target text;

102

accessing the source text from said first buffer and transducing the source text into at least one intermediate source structure of text constrained by any of said user defined criteria;

generating at least one target hypothesis, each of said target hypotheses comprising an intermediate target structure of text selected from the second language constrained by any of said user defined criteria;

estimating a first score, said first score being proportional to a probability of occurrence of each intermediate target structure of text associated with said target hypotheses using a target structure language model;

estimating a second score, said second score being proportional to a probability that said intermediate target structure of text associated with said target hypotheses will translate into said intermediate source structure of text using a target structure-to-source structure translation model;

combining, for each target hypothesis, said first and second scores to produce a target hypothesis match score;

transducing each of said intermediate target structures of text into at least one transformed target hypothesis of text in the second language constrained by any of said user defined criteria; and

performing at least one of a storing operation to save said at least one transformed target hypothesis in a second memory buffer and a presenting operation to make available for at least one of a viewing and listening operation according to its associated match score and said user defined criteria.

11. A method according to claim 10, wherein said intermediate source and target structures are transduced by:

arranging and rearranging, respectively, elements of the source and target text according to at least one of lexical substitutions, part-of-speech assignment, a morphological analysis, and a syntactic analysis.

12. A method according to claim 11, wherein said syntactic analysis further comprises the steps of:

performing question inversion;

performing do-not coalescence; and

performing adverb movement.

13. A method according to claim 11, wherein said intermediate source structure transducing step further comprises the step of:

annotating said words with parts of speech by probabilistically assigning a parts of speech to said words according to a statistical model.

14. A method according to claim 13, wherein said intermediate source structure transducing step further comprises the step of:

analyzing, syntactically, said annotated words and outputting a sequence of words and a sequence of said parts of speech corresponding to said sequence of words.

15. A method according to claim 14, wherein said intermediate source structure transducing step further comprises the step of:

analyzing, morphologically, said syntactically analyzed words by performing a morphological analysis of said words in accordance with said assigned parts of speech.

16. A method according to claim 15, wherein said intermediate source structure transducing step further comprises the step of:

assigning a sense to at least one of said morphological units to elucidate the translation of that morphological unit in the target language.

103

17. A method according to claim 11, wherein said syntactic analysis further comprises the steps of:

- performing question inversion;
- performing negative coalescence; and
- performing at least one of adverb, pronoun and adjective movement.

18. A method according to claim 10, wherein said intermediate source structure transducing step further comprises the step of:

- tokenizing the source text by identifying individual words and word separators and arranging said words and said word separators in a sequence.

19. A method according to claim 18, wherein said intermediate source structure transducing step further comprises the step of:

- determining case of said words using a case transducer, said case transducer assigning each word a token and a case pattern, said case pattern specifying a case of each letter of the word, and evaluating each case pattern to determine a true case pattern of the word.

20. A method according to claim 19, wherein said determining step further comprises the steps of:

- (1) determining whether said token is part of a name;
- (2) if said token is not part of a name:
  - determining whether said token is a member of a set of tokens which have only one true case; if said token is part of a name:
  - setting the true case pattern according to a table of true cases for names;
- (3) if the true case has not been determined by steps 1 and 2, and if the token is a first token in a sentence, then:
  - setting the true case pattern equal to a most probable true case for that token; and
- (4) otherwise:
  - setting the true case pattern equal to the only one case for that token.

21. A method according to claim 10, wherein each of said intermediate target structures is expressed as an ordered sequence of units of linguistic structure, and a first score probability is obtained by multiplying conditional probabilities of each of said units of linguistic structure within an intermediate target structure given an occurrence of previous units of linguistic structure within said intermediate target structure.

22. A method according to claim 21, wherein said conditional probability of each unit of each of linguistic structure within an intermediate target structures depends only on a fixed number of preceding units within said intermediate target structure.

23. A computer system for translating source text from a first language to target text in a second language different from the first language, said system comprising:

- means for receiving the source text in the first language;
- means for generating at least one target hypothesis, each of said target hypotheses comprising text selected from the second language;

- means for estimating, for each target hypothesis, a first probability of occurrence of said text associated with said target hypothesis using a target language model;

- means for estimating, for each target hypothesis, a second probability of occurrence of the source text given the occurrence of said text associated with said target hypothesis using a target-to-source translation model;

- means for combining, for each target hypothesis, said first and second probabilities to produce a target hypothesis match score; and

104

- means for performing at least one of a storing and presenting operation to save or otherwise make available for at least one of a viewing and listening operation, at least one of said target hypotheses according to its associated match score.

24. A system according to claim 23, further comprising: means for generating a list of partial target hypotheses using said target language model and said target-to-source translation model;

- means for extending said partial hypotheses to generate a set of new hypotheses;

- means for determining a match score for at least one of said new hypotheses; and

- means for performing at least one of a storing and presenting operation to save or otherwise make available for at least one of a viewing and listening operation, at least one of said new hypotheses according to its associated match score.

25. A system according to claim 24, wherein said means for determining said match score further comprises:

- means for determining a language model score for each of said partial hypotheses using said target language model;

- means for determining a translation model score for each of said partial hypotheses using said target-to-source translation model;

- and means for combining said language and said translation model scores for each of said partial hypotheses into a combined score.

26. A system according to claim 25, wherein said means for determining said translation model score further comprises:

- means for determining a fertility model score for each of said partial hypotheses, said partial hypotheses comprising at least one notion and said fertility model score being proportional to the probability that a notion in the target text will generate a specific number of notion units of linguistic structure in an intermediate structure source text;

- means for determining an alignment score for each of said partial hypotheses, said alignment score being proportional to a probability that a unit of linguistic structure in said intermediate structure of target text will align with one of zero or more units of linguistic structure in said intermediate structure of source text;

- means for determining a lexical model score, for each of said partial hypotheses, said lexical model score being proportional to the probability that said units of linguistic structure in said intermediate structure of target text of a given partial hypothesis will translate into said units of linguistic structure of said intermediate structure source text;

- means for determining a distortion model score for each of said partial hypotheses, said distortion model score being proportional to a probability that a source unit of linguistic structure will be in a particular position given the position of the target units of linguistic structure that generated it; and

- means for determining a combined score, said combined score being proportional to the product of at least one of said fertility, alignment, lexical and distortion model scores for each of said partial hypotheses.

27. A system according to claim 25, wherein said means for determining said translation model score further comprises:

means for determining a fertility model score for each of said partial hypotheses, said partial hypotheses comprising at least one notion and said fertility model score being proportional to the probability that a notion in the target text will generate a specific number of notions of units of linguistic structure in the source text; 5

means for determining an alignment score for each of said partial hypotheses, said alignment score being proportional to the probability that a unit of linguistic structure in the target text will align with one of zero or more units of linguistic structure in the source text; 10

means for determining a lexical model score, for each of said partial hypotheses, said lexical model score being proportional to the probability that said units of linguistic structure in the target text of a given partial hypothesis will translate into said units of linguistic structure of source text; 15

means for determining a distortion model score for each of said partial hypothesis, said distortion model score being proportional to the probability that the source units of linguistic structure will be in a particular position given the position of the target units of linguistic structure that generated it; and 20

means for determining a combined score, said combined score being proportional to the product of at least one of said fertility, alignment, lexical and distortion model scores for each of said partial hypotheses. 25

28. A system according to claim 27, wherein each of said notions comprises at least one unit of linguistic structure, and each of said units of linguistic structure in the target text produces one of zero or more units of linguistic structure in the source text. 30

29. A system according to claim 23, further comprising: means for receiving at least one user defined criteria pertaining to the source text to thereby bound the target text, said criteria belonging to a group comprising specification of a noun phrase, specification of subject, specification of verb, a semantic sense of a word or group of words, inclusion of a predetermined word or group of words and exclusion of a predetermined word or group of words. 35

30. A computer system for translating source text from a first language to target text in a second language different from the first language, said system comprising: 40

means for receiving the source text in the first language and storing the source text in a first memory buffer;

means for receiving one of zero or more user defined criteria pertaining to the source and target texts to thereby bound the target text; 45

means for accessing the source text from said first buffer; first transducing means for transducing the source text into at least one intermediate source structure of text constrained by any of said user defined criteria; 50

means for generating at least one target hypothesis, each of said target hypotheses comprising an intermediate target structure of text selected from the second language constrained by any of said user defined criteria; 55

means for estimating a first score, said first score being proportional to a probability of occurrence of each intermediate target structure of text associated with said target hypotheses using a target structure language model; 60

means for estimating a second score, said second score being proportional to a probability that said intermediate target structure of text associated with said target 65

hypotheses will translate into said intermediate source structure of text using a target structure-to-source structure translation model;

means for combining, for each target hypothesis, said first and second scores to produce a target hypothesis match score;

second transducing means for transducing each of said intermediate target structures of text into at least one transformed target hypothesis of text in the second language constrained by any of said user defined criteria; and

means for at least one of storing said at least one transformed target hypothesis in a second memory buffer, and presenting or otherwise making said at least one transformed target hypothesis available for at least one of a viewing and listening operation according to its associated match score and said user defined criteria.

31. A system according to claim 30, wherein said first and second transducing means further comprise:

third means for arranging elements of the source text according to at least one of lexical substitutions, part-of-speech assignment, a morphological analysis, and a syntactic analysis; and fourth means for rearranging elements of the target text according to at least one of lexical substitutions, part-of-speech assignment, a morphological analysis, and a syntactic analysis.

32. A system according to claim 31, further comprising: means for performing question inversion;

means for performing do-not coalescence; and

means for performing adverb movement.

33. A system according to claim 31, further comprising: means for annotating said words with parts of speech by probabilistically assigning a parts of speech of said words according to a statistical model.

34. A system according to claim 33, further comprising: means for analyzing, syntactically, said annotated words and outputting a sequence of words and a sequence of said part-of-speech corresponding to said sequence of words.

35. A system according to claim 34, further comprises: means for analyzing, morphologically, said syntactically analyzed words by performing a morphological analysis of said words in accordance with said assigned parts of speech.

36. A system according to claim 35, further comprising: means for assigning a sense to at least one of said morphological units to elucidate the translation of that morphological unit in the target language.

37. A system according to claim 31, further comprising: means for performing question inversion;

means for performing negative coalescence; and

means for performing at least one of adverb, pronoun and adjective movement.

38. A system according to claim 30, wherein said first transducing means further comprises;

means for tokenizing the source text by identifying individual units of linguistic structure and unit separators and arranging said units of linguistic structure and said unit separators in a sequence.

39. A system according to claim 38, further comprising: a case transducer means for assigning to each word a token and a case pattern, said case pattern specifying a case of each letter of the word, and evaluating each case pattern to determine a true case pattern of the word.

## 107

40. A system according to claim 39, further comprising:
- (a) means for determining whether said token is part of a name;
  - (b) means for setting a true case according to a table of true cases for names if said token is part of a name; 5
  - (c) means for determining whether said token is a member of a set of tokens which have only one true case if said token is not part of a name;
  - (c) if the true case has not be determined by steps a and b, and if said token is part of a name, then: 10
  - (d) means for setting the true case equal to a most probable true case for that token if said token is part of a name; and
  - (f) means for setting the true case equal to the case for that token if said token is not part of a name. 15

## 108

41. A system according to claim 30, further comprising:
- means for expressing each of said intermediate target structures as an ordered sequence of units of linguistic structure, and means for multiplying conditional probabilities of said units within an intermediate target structure given an occurrence of previous units of linguistic structure within said intermediate target structure to obtain said first score.
42. A system according to claim 41, wherein a conditional probability of each unit of each of said intermediate target structure depends only on a fixed number of preceding units within said intermediate target structure.

\* \* \* \* \*

UNITED STATES PATENT AND TRADEMARK OFFICE  
**CERTIFICATE OF CORRECTION**

PATENT NO. : 5,477,451

DATED : December 19, 1995

INVENTOR(S) : Peter F. Brown, et al.

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

Column 99:

Claim 1, line 5, change "firs" to --first--;

Column 100:

Claim 6, line 10, change "a" to --an--,

line 28 (column 101, line 5), change "leas tone" to --least one--; and

Signed and Sealed this  
Eighteenth Day of June, 1996

Attest:



BRUCE LEHMAN

Attesting Officer

Commissioner of Patents and Trademarks